

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
29 March 2001 (29.03.2001)


PCT

(10) International Publication Number
WO 01/22729 A1

- (51) International Patent Classification⁷: **H04N 5/92**
- (21) International Application Number: **PCT/US00/25847**
- (22) International Filing Date:
20 September 2000 (20.09.2000)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/154,713 20 September 1999 (20.09.1999) US
Not furnished 20 September 2000 (20.09.2000) US
- (71) Applicant: **TIVO, INC.** [US/US]; 2160 Gold Street, P.O. Box 2160, Alviso, CA 95002 (US).
- (72) Inventors: **BARTON, James, M.**; 101 Sund Avenue, Los Gatos, CA 95032 (US). **SMITH, Kevin**; 1164 Karen Way, Mountain View, CA 94040 (US). **CHAMBERLIN, David**; 206 Flynn Avenue, Mountain View, CA 94043 (US). **LOOK, Howard**; 576 Palo Alto Avenue, Mountain View, CA 94041 (US).
- (74) Agents: **GLENN, Michael et al.**; Glenn Patent Group, 3475 Edison Way, Ste. L., Menlo Park, CA 94025 (US).
- (81) Designated States (*national*): AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— *With international search report.*
— *Before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments.*

[Continued on next page]

(54) Title: **CLOSED CAPTION TAGGING SYSTEM**



(57) **Abstract:** A closed caption tagging system provides a mechanism for inserting tags into an audio or video television broadcast stream prior to or at the time of transmission. The tags contain command and control information that the receiver translates and acts upon. The receiver receives the broadcast stream and detects and processes the tags within the broadcast stream which is stored on a storage device that resides on the receiver. Program material from the broadcast stream is played back to the viewer from the storage device. The receiver performs the appropriate actions in response to the tags. Tags indicate the start and end points of a program segment. The receiver skips over a program segment during playback in response to the viewer pressing a button on a remote input device or it automatically skips over program segments depending on the viewer's preferences. Program segments such as commercials are automatically replaced by the receiver with new program segments that are selected based on various criteria. Menus, icons, and Web pages are displayed to the viewer based on information included in a tag. The viewer interacts with the menu, icon, or Web page through an input device with the receiver performing the associated actions. If a menu or actions requires that the viewer exits from the playback of the program material, then the receiver saves the exit point and returns the viewer back to the same exit point when the viewer has completed the interaction session. Menus and icons are used to generate leads, generate sales, and schedule the recording of programs. A one-touch recording option is provided. An icon is displayed to the viewer telling the viewer that an advertised program is available for recording at a future time. The viewer presses a single button on an input device causing the receiver to schedule the program for recording.

WO 01/22729 A1



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

CLOSED CAPTION TAGGING SYSTEM

5

BACKGROUND OF THE INVENTION

TECHNICAL FIELD

10

The invention relates to the processing of multimedia audio and video streams. More particularly, the invention relates to the tagging of multimedia audio and video television streams.

DESCRIPTION OF THE PRIOR ART

15

The Video Cassette Recorder (VCR) has changed the lives of television (TV) viewers throughout the world. The VCR has offered viewers the flexibility to time-shift TV programs to match their lifestyles.

20

The viewer stores TV programs onto magnetic tape using the VCR. The VCR gives the viewer the ability to play, rewind, fast forward and pause the stored program material. These functions enable the viewer to pause the program playback whenever he desires, fast forward through unwanted program material or commercials, and to replay favorite scenes. However, a VCR cannot both capture and play back information at the same time.

25

Digital Video Recorders (DVR) have recently entered into the marketplace. DVRs allow the viewer to store TV programs on a hard disk. This has freed the viewer from the magnetic tape realm. Viewers can pause, rewind, and fast forward live broadcast programs. However, the functionality of DVRs extends beyond recording programs.

30

35

Having programs stored locally in a digital form gives the programmer many more options than were previously available. Advertisements (commercials) can now be dynamically replaced and specifically targeted to the particular viewer based on his or her viewing habits. The commercials can be stored locally on the viewer's DVR and shown at any time.

DVRs allow interactive programming with the viewer. Generally, promotions for future shows are displayed to viewers during the normal broadcast programs. Viewers must then remember the date, time, and channel that the program will be aired on to record or view the program. DVRs allow the viewer to schedule the recording of the program immediately.

The only drawback is that the current generation of DVRs do not have the capability to interact with the viewer at this level. There is no means by which to notify the DVR that commercials are directly tied to a certain program or other advertisements. Further, there is no way to tell the DVR that a commercial can be replaced.

It would be advantageous to provide a closed caption tagging system that gives the content provider the ability to send frame specific data across broadcast media. It would further be advantageous to provide a closed caption tagging system that allows the receiver to dynamically interact with the viewer and configure itself based on program content.

SUMMARY OF THE INVENTION

The invention provides a closed caption tagging system. The invention allows content providers to send frame specific data and commands integrated into video and audio television streams across broadcast media. In addition, the invention allows the receiver to dynamically interact with the viewer and configure itself based on video and audio stream content.

A preferred embodiment of the invention provides a mechanism for inserting tags into an audio or video television broadcast stream. Tags are inserted into the broadcast stream prior to or at the time of transmission. The tags contain command and control information that the receiver translates and acts upon.

The receiver receives the broadcast stream and detects and processes the tags within the broadcast stream. The broadcast stream is stored on a storage device that resides on the receiver. Program material from the broadcast stream is played back to the viewer from the storage device.

During the tag processing stage, the receiver performs the appropriate actions in

response to the tags. The tags offer a great amount of flexibility to the content provider or system administrator to create a limitless amount of operations.

5 Tags indicate the start and end points of a program segment. The receiver skips over a program segment during playback in response to the viewer pressing a button on a remote input device. The receiver also automatically skips over program segments depending on the viewer's preferences.

10 Program segments such as commercials are automatically replaced by the receiver with new program segments. New program segments are selected based on various criteria such as the locale, time of day, program material, viewer's viewing habits, viewer's program preferences, or the viewer's personal information. The new program segments are stored remotely or locally on the receiver.

15 Menus, icons, and Web pages are displayed to the viewer based on information included in a tag. The viewer interacts with the menu, icon, or Web page through an input device. The receiver performs the actions associated with the menu, icon, or Web page and the viewer's input. If a menu or action requires
20 that the viewer exit from the playback of the program material, then the receiver saves the exit point and returns the viewer back to the same exit point when the viewer has completed the interaction session.

25 Menus and icons are used to generate leads, generate sales, and schedule the recording of programs. A one-touch recording option is provided. An icon is displayed to the viewer telling the viewer that an advertised program is available for recording at a future time. The viewer presses a single button on an input device causing the receiver to schedule the program for recording. The receiver will also record the current program in the broadcast stream onto the storage
30 device based on information included in a tag.

Tags are used to create indexes in program material. This allows the viewer to jump to particular indexes in a program.

35 Other aspects and advantages of the invention will become apparent from the following detailed description in combination with the accompanying drawings, illustrating, by way of example, the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

- 5 Fig. 1 is a block schematic diagram of a high level view of a preferred embodiment of the invention according to the invention;
- Fig. 2 is a block schematic diagram of a preferred embodiment of the invention using multiple input and output modules according to the invention;
- 10 Fig. 3 is a schematic diagram of an Moving Pictures Experts Group (MPEG) data stream and its video and audio components according to the invention;
- Fig. 4 is a block schematic diagram of a parser and four direct memory access (DMA) input engines contained in the Media Switch according to the invention;
- 15 Fig. 5 is a schematic diagram of the components of a packetized elementary stream (PES) buffer according to the invention;
- Fig. 6 is a schematic diagram of the construction of a PES buffer from the parsed components in the Media Switch output circular buffers;
- 20 Fig. 7 is a block schematic diagram of the Media Switch and the various components that it communicates with according to the invention;
- 25 Fig. 8 is a block schematic diagram of a high level view of the program logic according to the invention;
- Fig. 9 is a block schematic diagram of a class hierarchy of the program logic according to the invention;
- 30 Fig. 10 is a block schematic diagram of a preferred embodiment of the clip cache component of the invention according to the invention;
- 35 Fig. 11 is a block schematic diagram of a preferred embodiment of the invention that emulates a broadcast studio video mixer according to the invention;

Fig. 12 is a block schematic diagram of a closed caption parser according to the invention;

Fig. 13 is a block schematic diagram of a high level view of a preferred embodiment of the invention utilizing a VCR as an integral component of the invention according to the invention;

Fig. 14 is a block schematic diagram of a preferred embodiment of the invention for inserting tags into a video stream according to the invention;

Fig. 15 is a block schematic diagram of a server-based preferred embodiment of the invention for inserting tags into a video stream according to the invention;

Fig. 16 is a diagram of a user interface for inserting tags into a video stream according to the invention;

Fig. 17 is a diagram of a screen with an alert icon displayed in the lower left corner of the screen according to the invention;

Fig. 18 is a block schematic diagram of the transmission route of a video stream according to the invention;

Fig. 19 is a block schematic diagram of the tagging of the start and end of a program segment of a video stream and the playback of a new program segment according to the invention;

Fig. 20 is a block schematic diagram of a preferred embodiment of the invention that interprets tags inserted into a video stream according to the invention;

Fig. 21 is a diagram of a screen displaying program recording options according to the invention;

Fig. 22 is a diagram of a viewer remote control device according to the invention; and

Fig. 23 is a block schematic diagram of a series of screens for lead and sale generation according to the invention.

DETAILED DESCRIPTION OF THE INVENTION

5 The invention is embodied in a closed caption tagging system. A system according to the invention allows content providers to send frame specific data and commands integrated into video and audio television streams across broadcast media. The invention additionally allows the receiver to dynamically interact with the viewer and configure itself based on video and audio stream content.

10 A preferred embodiment of the invention provides a tagging and interpretation system that allows a content provider to tag, in a frame specific manner, video and audio streams transmitted over television broadcast media. A receiver interprets and acts upon the tags embedded in the received stream. The tag
15 data allow the receiver to dynamically interact with the viewer through menus and action icons. The tags also provide for the dynamic configuration of the receiver.

Referring to Fig. 1, a preferred embodiment of the invention has an Input Section
20 101, Media Switch 102, and an Output Section 103. The Input Section 101 takes television (TV) input streams in a multitude of forms, for example, National Television Standards Committee (NTSC) or PAL broadcast, and digital forms such as Digital Satellite System (DSS), Digital Broadcast Services (DBS), or Advanced Television Standards Committee (ATSC). DBS, DSS and ATSC are based on standards called Moving Pictures Experts Group 2 (MPEG2) and
25 MPEG2 Transport. MPEG2 Transport is a standard for formatting the digital data stream from the TV source transmitter so that a TV receiver can disassemble the input stream to find programs in the multiplexed signal. The Input Section 101 produces MPEG streams. An MPEG2 transport multiplex supports multiple programs in the same broadcast channel, with multiple video and audio feeds and private data. The Input Section 101 tunes the channel to a particular
30 program, extracts a specific MPEG program out of it, and feeds it to the rest of the system. Analog TV signals are encoded into a similar MPEG format using separate video and audio encoders, such that the remainder of the system is unaware of how the signal was obtained. Information may be modulated into the Vertical Blanking Interval (VBI) of the analog TV signal in a number of standard
35 ways; for example, the North American Broadcast Teletext Standard (NABTS) may be used to modulate information onto lines 10 through 20 of an NTSC

signal, while the FCC mandates the use of line 21 for Closed Caption (CC) and Extended Data Services (EDS). Such signals are decoded by the input section and passed to the other sections as if they were delivered via an MPEG2 private data channel.

5

10 The Media Switch 102 mediates between a microprocessor CPU 106, hard disk or storage device 105, and memory 104. Input streams are converted to an MPEG stream and sent to the Media Switch 102. The Media Switch 102 buffers the MPEG stream into memory. It then performs two operations if the user is watching real time TV: the stream is sent to the Output Section 103 and it is written simultaneously to the hard disk or storage device 105.

15 The Output Section 103 takes MPEG streams as input and produces an analog TV signal according to the NTSC, PAL, or other required TV standards. The Output Section 103 contains an MPEG decoder, On-Screen Display (OSD) generator, analog TV encoder and audio logic. The OSD generator allows the program logic to supply images which will be overlayed on top of the resulting analog TV signal. Additionally, the Output Section can modulate information supplied by the program logic onto the VBI of the output signal in a number of
20 standard formats, including NABTS, CC and EDS.

25 With respect to Fig. 2, the invention easily expands to accommodate multiple Input Sections (tuners) 201, 202, 203, 204, each can be tuned to different types of input. Multiple Output Modules (decoders) 206, 207, 208, 209 are added as well. Special effects such as picture in a picture can be implemented with multiple decoders. The Media Switch 205 records one program while the user is watching another. This means that a stream can be extracted off the disk while another stream is being stored onto the disk.

30 Referring to Fig. 3, the incoming MPEG stream 301 has interleaved video 302, 305, 306 and audio 303, 304, 307 segments. These elements must be separated and recombined to create separate video 308 and audio 309 streams or buffers. This is necessary because separate decoders are used to convert MPEG elements back into audio or video analog components. Such separate
35 delivery requires that time sequence information be generated so that the decoders may be properly synchronized for accurate playback of the signal.

The Media Switch enables the program logic to associate proper time sequence

information with each segment, possibly embedding it directly into the stream. The time sequence information for each segment is called a time stamp. These time stamps are monotonically increasing and start at zero each time the system boots up. This allows the invention to find any particular spot in any particular video segment. For example, if the system needs to read five seconds into an incoming contiguous video stream that is being cached, the system simply has to start reading forward into the stream and look for the appropriate time stamp.

A binary search can be performed on a stored file to index into a stream. Each stream is stored as a sequence of fixed-size segments enabling fast binary searches because of the uniform timestamping. If the user wants to start in the middle of the program, the system performs a binary search of the stored segments until it finds the appropriate spot, obtaining the desired results with a minimal amount of information. If the signal were instead stored as an MPEG stream, it would be necessary to linearly parse the stream from the beginning to find the desired location.

With respect to Fig. 4, the Media Switch contains four input Direct Memory Access (DMA) engines 402, 403, 404, 405 each DMA engine has an associated buffer 410, 411, 412, 413. Conceptually, each DMA engine has a pointer 406, a limit for that pointer 407, a next pointer 408, and a limit for the next pointer 409. Each DMA engine is dedicated to a particular type of information, for example, video 402, audio 403, and parsed events 405. The buffers 410, 411, 412, 413 are circular and collect the specific information. The DMA engine increments the pointer 406 into the associated buffer until it reaches the limit 407 and then loads the next pointer 408 and limit 409. Setting the pointer 406 and next pointer 408 to the same value, along with the corresponding limit value creates a circular buffer. The next pointer 408 can be set to a different address to provide vector DMA.

The input stream flows through a parser 401. The parser 401 parses the stream looking for MPEG distinguished events indicating the start of video, audio or private data segments. For example, when the parser 401 finds a video event, it directs the stream to the video DMA engine 402. The parser 401 buffers up data and DMA's it into the video buffer 410 through the video DMA engine 402. At the same time, the parser 401 directs an event to the event DMA engine 405 which generates an event into the event buffer 413. When the parser 401 sees an audio event, it redirects the byte stream to the audio DMA engine 403 and

generates an event into the event buffer 413. Similarly, when the parser 401 sees a private data event, it directs the byte stream to the private data DMA engine 404 and directs an event to the event buffer 413. The Media Switch notifies the program logic via an interrupt mechanism when events are placed in the event buffer.

Referring to Figs. 4 and 5, the event buffer 413 is filled by the parser 401 with events. Each event 501 in the event buffer has an offset 502, event type 503, and time stamp field 504. The parser 401 provides the type and offset of each event as it is placed into the buffer. For example, when an audio event occurs, the event type field is set to an audio event and the offset indicates the location in the audio buffer 411. The program logic knows where the audio buffer 411 starts and adds the offset to find the event in the stream. The address offset 502 tells the program logic where the next event occurred, but not where it ended. The previous event is cached so the end of the current event can be found as well as the length of the segment.

With respect to Figs. 5 and 6, the program logic reads accumulated events in the event buffer 602 when it is interrupted by the Media Switch 601. From these events the program logic generates a sequence of logical segments 603 which correspond to the parsed MPEG segments 615. The program logic converts the offset 502 into the actual address 610 of each segment, and records the event length 609 using the last cached event. If the stream was produced by encoding an analog signal, it will not contain Program Time Stamp (PTS) values, which are used by the decoders to properly present the resulting output. Thus, the program logic uses the generated time stamp 504 to calculate a simulated PTS for each segment and places that into the logical segment timestamp 607. In the case of a digital TV stream, PTS values are already encoded in the stream. The program logic extracts this information and places it in the logical segment timestamp 607.

The program logic continues collecting logical segments 603 until it reaches the fixed buffer size. When this occurs, the program logic generates a new buffer, called a Packetized Elementary Stream (PES) 605 buffer containing these logical segments 603 in order, plus ancillary control information. Each logical segment points 604 directly to the circular buffer, e.g., the video buffer 613, filled by the Media Switch 601. This new buffer is then passed to other logic components, which may further process the stream in the buffer in some way, such as

presenting it for decoding or writing it to the storage media. Thus, the MPEG data is not copied from one location in memory to another by the processor. This results in a more cost effective design since lower memory bandwidth and processor bandwidth is required.

5

10

15

A unique feature of the MPEG stream transformation into PES buffers is that the data associated with logical segments need not be present in the buffer itself, as presented above. When a PES buffer is written to storage, these logical segments are written to the storage medium in the logical order in which they appear. This has the effect of gathering components of the stream, whether they be in the video, audio or private data circular buffers, into a single linear buffer of stream data on the storage medium. The buffer is read back from the storage medium with a single transfer from the storage media, and the logical segment information is updated to correspond with the actual locations in the buffer 606. Higher level program logic is unaware of this transformation, since it handles only the logical segments, thus stream data is easily managed without requiring that the data ever be copied between locations in DRAM by the CPU.

20

25

A unique aspect of the Media Switch is the ability to handle high data rates effectively and inexpensively. It performs the functions of taking video and audio data in, sending video and audio data out, sending video and audio data to disk, and extracting video and audio data from the disk on a low cost platform. Generally, the Media Switch runs asynchronously and autonomously with the microprocessor CPU, using its DMA capabilities to move large quantities of information with minimal intervention by the CPU.

30

35

Referring to Fig. 7, the input side of the Media Switch 701 is connected to an MPEG encoder 703. There are also circuits specific to MPEG audio 704 and vertical blanking interval (VBI) data 702 feeding into the Media Switch 701. If a digital TV signal is being processed instead, the MPEG encoder 703 is replaced with an MPEG2 Transport Demultiplexor, and the MPEG audio encoder 704 and VBI decoder 702 are deleted. The demultiplexor multiplexes the extracted audio, video and private data channel streams through the video input Media Switch port.

The parser 705 parses the input data stream from the MPEG encoder 703, audio encoder 704 and VBI decoder 702, or from the transport demultiplexor in the case of a digital TV stream. The parser 705 detects the beginning of all of the

important events in a video or audio stream, the start of all of the frames, the start of sequence headers - all of the pieces of information that the program logic needs to know about in order to both properly play back and perform special effects on the stream, *e.g.* fast forward, reverse, play, pause, fast/slow play, indexing, and fast/slow reverse play.

The parser 705 places tags 707 into the FIFO 706 when it identifies video or audio segments, or is given private data. The DMA 709 controls when these tags are taken out. The tags 707 and the DMA addresses of the segments are placed into the event queue 708. The frame type information, whether it is a start of a video I-frame, video B-frame, video P-frame, video PES, audio PES, a sequence header, an audio frame, or private data packet, is placed into the event queue 708 along with the offset in the related circular buffer where the piece of information was placed. The program logic operating in the CPU 713 examines events in the circular buffer after it is transferred to the DRAM 714.

The Media Switch 701 has a data bus 711 that connects to the CPU 713 and DRAM 714. An address bus 712 is also shared between the Media Switch 701, CPU 713, and DRAM 714. A hard disk or storage device 710 is connected to one of the ports of the Media Switch 701. The Media Switch 701 outputs streams to an MPEG video decoder 715 and a separate audio decoder 717. The audio decoder 717 signals contain audio cues generated by the system in response to the user's commands on a remote control or other internal events. The decoded audio output from the MPEG decoder is digitally mixed 718 with the separate audio signal. The resulting signals contain video, audio, and on-screen displays and are sent to the TV 716.

The Media Switch 701 takes in 8-bit data and sends it to the disk, while at the same time extracts another stream of data off of the disk and sends it to the MPEG decoder 715. All of the DMA engines described above can be working at the same time. The Media Switch 701 can be implemented in hardware using a Field Programmable Gate Array (FPGA), ASIC, or discrete logic.

Rather than having to parse through an immense data stream looking for the start of where each frame would be, the program logic only has to look at the circular event buffer in DRAM 714 and it can tell where the start of each frame is and the frame type. This approach saves a large amount of CPU power, keeping the real time requirements of the CPU 713 small. The CPU 713 does not have to

be very fast at any point in time. The Media Switch 701 gives the CPU 713 as much time as possible to complete tasks. The parsing mechanism 705 and event queue 708 decouple the CPU 713 from parsing the audio, video, and buffers and the real time nature of the streams, which allows for lower costs. It also allows the use of a bus structure in a CPU environment that operates at a much lower clock rate with much cheaper memory than would be required otherwise.

The CPU 713 has the ability to queue up one DMA transfer and can set up the next DMA transfer at its leisure. This gives the CPU 713 large time intervals within which it can service the DMA controller 709. The CPU 713 may respond to a DMA interrupt within a larger time window because of the large latency allowed. MPEG streams, whether extracted from an MPEG2 Transport or encoded from an analog TV signal, are typically encoded using a technique called Variable Bit Rate encoding (VBR). This technique varies the amount of data required to represent a sequence of images by the amount of movement between those images. This technique can greatly reduce the required bandwidth for a signal, however sequences with rapid movement (such as a basketball game) may be encoded with much greater bandwidth requirements. For example, the Hughes DirecTV satellite system encodes signals with anywhere from 1 to 10Mb/s of required bandwidth, varying from frame to frame. It would be difficult for any computer system to keep up with such rapidly varying data rates without this structure.

With respect to Fig. 8, the program logic within the CPU has three conceptual components: sources 801, transforms 802, and sinks 803. The sources 801 produce buffers of data. Transforms 802 process buffers of data and sinks 803 consume buffers of data. A transform is responsible for allocating and queuing the buffers of data on which it will operate. Buffers are allocated as if "empty" to sources of data, which give them back "full". The buffers are then queued and given to sinks as "full", and the sink will return the buffer "empty".

A source 801 accepts data from encoders, *e.g.*, a digital satellite receiver. It acquires buffers for this data from the downstream transform, packages the data into a buffer, then pushes the buffer down the pipeline as described above. The source object 801 does not know anything about the rest of the system. The sink 803 consumes buffers, taking a buffer from the upstream transform, sending the data to the decoder, and then releasing the buffer for reuse.

There are two types of transforms 802 used: spatial and temporal. Spatial transforms are transforms that perform, for example, an image convolution or compression/decompression on the buffered data that is passing through.

Temporal transforms are used when there is no time relation that is expressible between buffers going in and buffers coming out of a system. Such a transform writes the buffer to a file 804 on the storage medium. The buffer is pulled out at a later time, sent down the pipeline, and properly sequenced within the stream.

Referring to Fig. 9, a C++ class hierarchy derivation of the program logic is shown. The TiVo Media Kernel (Tmk) 904, 908, 913 mediates with the operating system kernel. The kernel provides operations such as: memory allocation, synchronization, and threading. The TmkCore 904, 908, 913 structures memory taken from the media kernel as an object. It provides operators, new and delete, for constructing and deconstructing the object. Each object (source 901, transform 902, and sink 903) is multi-threaded by definition and can run in parallel.

The TmkPipeline class 905, 909, 914 is responsible for flow control through the system. The pipelines point to the next pipeline in the flow from source 901 to sink 903. To pause the pipeline, for example, an event called "pause" is sent to the first object in the pipeline. The event is relayed on to the next object and so on down the pipeline. This all happens asynchronously to the data going through the pipeline. Thus, similar to applications such as telephony, control of the flow of MPEG streams is asynchronous and separate from the streams themselves. This allows for a simple logic design that is at the same time powerful enough to support the features described previously, including pause, rewind, fast forward and others. In addition, this structure allows fast and efficient switching between stream sources, since buffered data can be simply discarded and decoders reset using a single event, after which data from the new stream will pass down the pipeline. Such a capability is needed, for example, when switching the channel being captured by the input section, or when switching between a live signal from the input section and a stored stream.

The source object 901 is a TmkSource 906 and the transform object 902 is a TmkXfrm 910. These are intermediate classes that define standard behaviors for the classes in the pipeline. Conceptually, they handshake buffers down the pipeline. The source object 901 takes data out of a physical data source, such as

the Media Switch, and places it into a PES buffer. To obtain the buffer, the source object 901 asks the down stream object in his pipeline for a buffer (allocEmptyBuf). The source object 901 is blocked until there is sufficient memory. This means that the pipeline is self-regulating; it has automatic flow control. When the source object 901 has filled up the buffer, it hands it back to the transform 902 through the pushFullBuf function.

The sink 903 is flow controlled as well. It calls nextFullBuf which tells the transform 902 that it is ready for the next filled buffer. This operation can block the sink 903 until a buffer is ready. When the sink 903 is finished with a buffer (*i.e.*, it has consumed the data in the buffer) it calls releaseEmptyBuf. ReleaseEmptyBuf gives the buffer back to the transform 902. The transform 902 can then hand that buffer, for example, back to the source object 901 to fill up again. In addition to the automatic flow-control benefit of this method, it also provides for limiting the amount of memory dedicated to buffers by allowing enforcement of a fixed allocation of buffers by a transform. This is an important feature in achieving a cost-effective limited DRAM environment.

The MediaSwitch class 909 calls the allocEmptyBuf method of the TmkClipCache 912 object and receives a PES buffer from it. It then goes out to the circular buffers in the Media Switch hardware and generates PES buffers. The MediaSwitch class 909 fills the buffer up and pushes it back to the TmkClipCache 912 object.

The TmkClipCache 912 maintains a cache file 918 on a storage medium. It also maintains two pointers into this cache: a push pointer 919 that shows where the next buffer coming from the source 901 is inserted; and a current pointer 920 which points to the current buffer used.

The buffer that is pointed to by the current pointer is handed to the Vela decoder class 916. The Vela decoder class 916 talks to the decoder 921 in the hardware. The decoder 921 produces a decoded TV signal that is subsequently encoded into an analog TV signal in NTSC, PAL or other analog format. When the Vela decoder class 916 is finished with the buffer it calls releaseEmptyBuf.

The structure of the classes makes the system easy to test and debug. Each level can be tested separately to make sure it performs in the appropriate manner, and the classes may be gradually aggregated to achieve the desired

functionality while retaining the ability to effectively test each object.

The control object 917 accepts commands from the user and sends events into the pipeline to control what the pipeline is doing. For example, if the user has a remote control and is watching TV, the user presses pause and the control object 917 sends an event to the sink 903, that tells it pause. The sink 903 stops asking for new buffers. The current pointer 920 stays where it is at. The sink 903 starts taking buffers out again when it receives another event that tells it to play. The system is in perfect synchronization; it starts from the frame that it stopped at.

The remote control may also have a fast forward key. When the fast forward key is pressed, the control object 917 sends an event to the transform 902, that tells it to move forward two seconds. The transform 902 finds that the two second time span requires it to move forward three buffers. It then issues a reset event to the downstream pipeline, so that any queued data or state that may be present in the hardware decoders is flushed. This is a critical step, since the structure of MPEG streams requires maintenance of state across multiple frames of data, and that state will be rendered invalid by repositioning the pointer. It then moves the current pointer 920 forward three buffers. The next time the sink 903 calls nextFullBuf it gets the new current buffer. The same method works for fast reverse in that the transform 902 moves the current pointer 920 backwards.

A system clock reference resides in the decoder. The system clock reference is sped up for fast play or slowed down for slow play. The sink simply asks for full buffers faster or slower, depending on the clock speed.

With respect to Fig. 10, two other objects derived from the TmkXfrm class are placed in the pipeline for disk access. One is called TmkClipReader 1003 and the other is called TmkClipWriter 1001. Buffers come into the TmkClipWriter 1001 and are pushed to a file on a storage medium 1004. TmkClipReader 1003 asks for buffers which are taken off of a file on a storage medium 1005. A TmkClipReader 1003 provides only the allocEmptyBuf and pushFullBuf methods, while a TmkClipWriter 1001 provides only the nextFullBuf and releaseEmptyBuf methods. A TmkClipReader 1003 therefore performs the same function as the input, or "push" side of a TmkClipCache 1002, while a TmkClipWriter 1001 therefore performs the same function as the output, or "pull" side of a TmkClipCache 1002.

Referring to Fig. 11, a preferred embodiment that accomplishes multiple functions is shown. A source 1101 has a TV signal input. The source sends data to a PushSwitch 1102 which is a transform derived from TmkXfm. The PushSwitch 1102 has multiple outputs that can be switched by the control object 1114. This means that one part of the pipeline can be stopped and another can be started at the users whim. The user can switch to different storage devices. The PushSwitch 1102 could output to a TmkClipWriter 1106, which goes onto a storage device 1107 or write to the cache transform 1103.

An important feature of this apparatus is the ease with which it can selectively capture portions of an incoming signal under the control of program logic. Based on information such as the current time, or perhaps a specific time span, or perhaps via a remote control button press by the viewer, a TmkClipWriter 1106 may be switched on to record a portion of the signal, and switched off at some later time. This switching is typically caused by sending a "switch" event to the PushSwitch 1102 object.

An additional method for triggering selective capture is through information modulated into the VBI or placed into an MPEG private data channel. Data decoded from the VBI or private data channel is passed to the program logic. The program logic examines this data to determine if the data indicates that capture of the TV signal into which it was modulated should begin. Similarly, this information may also indicate when recording should end, or another data item may be modulated into the signal indicating when the capture should end. The starting and ending indicators may be explicitly modulated into the signal or other information that is placed into the signal in a standard fashion may be used to encode this information.

With respect to Fig. 12, an example is shown which demonstrates how the program logic scans the words contained within the closed caption (CC) fields to determine starting and ending times, using particular words or phrases to trigger the capture. A stream of NTSC or PAL fields 1201 is presented. CC bytes are extracted from each odd field 1202, and entered in a circular buffer 1203 for processing by the Word Parser 1204. The Word Parser 1204 collects characters until it encounters a word boundary, usually a space, period or other delineating character. Recall from above, that the MPEG audio and video segments are collected into a series of fixed-size PES buffers. A special segment is added to each PES buffer to hold the words extracted from the C C

field 1205. Thus, the CC information is preserved in time synchronization with the audio and video, and can be correctly presented to the viewer when the stream is displayed. This also allows the stored stream to be processed for CC information at the leisure of the program logic, which spreads out load, reducing cost and improving efficiency. In such a case, the words stored in the special segment are simply passed to the state table logic 1206.

During stream capture, each word is looked up in a table 1206 which indicates the action to take on recognizing that word. This action may simply change the state of the recognizer state machine 1207, or may cause the state machine 1207 to issue an action request, such as "start capture", "stop capture", "phrase seen", or other similar requests. Indeed, a recognized word or phrase may cause the pipeline to be switched; for example, to overlay a different audio track if undesirable language is used in the program.

Note that the parsing state table 1206 and recognizer state machine 1207 may be modified or changed at any time. For example, a different table and state machine may be provided for each input channel. Alternatively, these elements may be switched depending on the time of day, or because of other events.

Referring to Fig. 11, a PullSwitch is added 1104 which outputs to the sink 1105. The sink 1105 calls nextFullBuf and releaseEmptyBuf to get or return buffers from the PullSwitch 1104. The PullSwitch 1104 can have any number of inputs. One input could be an ActionClip 1113. The remote control can switch between input sources. The control object 1114 sends an event to the PullSwitch 1104, telling it to switch. It will switch from the current input source to whatever input source the control object selects.

An ActionClip class provides for sequencing a number of different stored signals in a predictable and controllable manner, possibly with the added control of viewer selection via a remote control. Thus, it appears as a derivative of a TmkXfrm object that accepts a "switch" event for switching to the next stored signal.

This allows the program logic or user to create custom sequences of video output. Any number of video segments can be lined up and combined as if the program logic or user were using a broadcast studio video mixer. TmkClipReaders 1108, 1109, 1110 are allocated and each is hooked into the

PullSwitch 1104. The PullSwitch 1104 switches between the TmkClipReaders 1108, 1109, 1110 to combine video and audio clips. Flow control is automatic because of the way the pipeline is constructed. The Push and Pull Switches are the same as video switches in a broadcast studio.

5

The derived class and resulting objects described here may be combined in an arbitrary way to create a number of different useful configurations for storing, retrieving, switching and viewing of TV streams. For example, if multiple input and output sections are available, one input is viewed while another is stored, and a picture-in-picture window generated by the second output is used to preview previously stored streams. Such configurations represent a unique and novel application of software transformations to achieve the functionality expected of expensive, sophisticated hardware solutions within a single cost-effective device.

10

15

With respect to Fig. 13, a high-level system view is shown which implements a VCR backup. The Output Module 1303 sends TV signals to the VCR 1307. This allows the user to record TV programs directly on to video tape. The invention allows the user to queue up programs from disk to be recorded on to video tape and to schedule the time that the programs are sent to the VCR 1307. Title pages (EPG data) can be sent to the VCR 1307 before a program is sent. Longer programs can be scaled to fit onto smaller video tapes by speeding up the play speed or dropping frames.

20

25

The VCR 1307 output can also be routed back into the Input Module 1301. In this configuration the VCR acts as a backup system for the Media Switch 1302. Any overflow storage or lower priority programming is sent to the VCR 1307 for later retrieval.

30

The Input Module 1301 can decode and pass to the remainder of the system information encoded on the Vertical Blanking Interval (VBI). The Output Module 1303 can encode into the output VBI data provided by the remainder of the system. The program logic may arrange to encode identifying information of various kinds into the output signal, which will be recorded onto tape using the VCR 1307. Playing this tape back into the input allows the program logic to read back this identifying information, such that the TV signal recorded on the tape is properly handled. For example, a particular program may be recorded to tape along with information about when it was recorded, the source network, etc.

35

When this program is played back into the Input Module, this information can be used to control storage of the signal, presentation to the viewer, etc.

One skilled in the art will readily appreciate that such a mechanism may be used to introduce various data items to the program logic which are not properly conceived of as television signals. For instance, software updates or other data may be passed to the system. The program logic receiving this data from the television stream may impose controls on how the data is handled, such as requiring certain authentication sequences and/or decrypting the embedded information according to some previously acquired key. Such a method works for normal broadcast signals as well, leading to an efficient means of providing non-TV control information and data to the program logic.

Additionally, one skilled in the art will readily appreciate that although a VCR is specifically mentioned above, any multimedia recording device (e.g., a Digital Video Disk-Random Access Memory (DVD-RAM) recorder) is easily substituted in its place.

Although the invention is described herein with reference to the preferred embodiment, one skilled in the art will readily appreciate that other applications may be substituted for those set forth herein without departing from the spirit and scope of the present invention. For example, the invention can be used in the detection of gambling casino crime. The input section of the invention is connected to the casino's video surveillance system. Recorded video is cached and simultaneously output to external VCRs. The user can switch to any video feed and examine (i.e., rewind, play, slow play, fast forward, etc.) a specific segment of the recorded video while the external VCRs are being loaded with the real-time input video.

Video Stream Tag Architecture

Referring again to Fig. 12, tags are abstract events which occur in a television stream 1201. They may be embedded in the VBI of an analog signal, or in a private data channel in an MPEG2 multiplex. As described above, tags can be embedded in the closed caption (CC) fields and extracted into a circular buffer 1203 or memory allocation schema. The word parser 1204 identifies unique tags during its scan of the CC data. Tags are interspersed with the standard CC control codes. Tags may also be generated implicitly, for instance, based on the

current time and program being viewed.

The invention provides a mechanism called the TiVo Video Tag Authoring (TVTAG) system for inserting tags (TiVo tags) into a video stream prior to broadcast. With respect to Figs. 14, 16, and 17, the TVTAG system consists of a video output source 1401, a compatible device for inserting Vertical Blanking Interval (VBI) closed-captioning information and outputting captioned video 1402, a video monitor 1405, and a software program for controlling the VBI insertion device to incorporate tag data objects in the form of closed-caption information in the video stream 1406. The tagged video is retransmitted immediately 1404 or stored on a suitable medium 1403 for later transmission.

The TVTAG software 1406, in its most basic implementation, is responsible for controlling the VBI Insertion device 1402. The TVTAG software 1406 communicates with the VBI insertion device 1402 by means of standard computer interfaces and device control code protocols. When an operator observing the video monitor 1405 determines that the desired tag insertion point has been reached, he presses a key, causing the TiVo tag data object to be generated, transmitted to the VBI insertion device 1402, and incorporated in the video stream for transmission 1404 or storage 1403.

The TVTAG software has the additional capability of controlling the video input source 1401 and the video output storage device 1403. The operator selects the particular video 1602 and has the ability to pause the video input stream to facilitate overlaying a graphic element 1702 on the monitor, and positioning it by means of a pointing device, such as a mouse. The positioning of the graphic element 1702 is also accomplished through the operator interface 1601. The operator inputs the position of the graphic using the X position 1605 and the Y position 1604.

The graphic element and positioning information are then incorporated in the TiVo tag data object (discussed below) and the time-code or frame of the video noted. When the operator is satisfied, playback and record are resumed. The tag is then issued through the insertion device with the highest degree of accuracy.

Referring to Fig. 15, in another preferred embodiment of the TVTAG system, the software program takes the form of a standard Internet protocol Web page

displayed to operator(s) 1505. The Web page causes the TiVo tag object to be generated by a script running on a remote server 1504. The server 1504 controls the VBI insertion device 1502, the video source 1501, and recording devices 1503. The remote operator(s) 1505 can receive from the server 1504 a
5 low or high-bandwidth version of the video stream for use as a reference for tag insertion. Once the necessary tag data object information has been generated and transmitted, it can be batch-processed at a later time by the server 1504.

Another preferred embodiment of the invention integrates the software with
10 popular non-linear video editing systems as a "plug-in", thereby allowing the TiVo tag data objects to be inserted during the video production process. In this embodiment, the non-linear editing system serves as the source and storage system controller and also provides graphic placement facilities, allowing frame-accurate placement of the TiVo tag data object.

15 With respect to Fig. 18, tags are integrated into the video stream before or at the video source 1801. The video stream is then transmitted via satellite 1802, cable or other terrestrial transmission method 1803. The receiver 1804 receives the video stream, recognizes the tags and performs the appropriate actions in
20 response to the tags. The viewer sees the resultant video stream via the monitor or television set 1805.

The invention provides an architecture that supports taking various actions based on tags in the video stream. Some examples of the flexibility that TiVo tags
25 offer are:

- It is desirable to know when a network promotion is being viewed so that the viewer might be presented with an option to record the program at some future time. TiVo tags are added into the promotion that indicate the date,
30 time, and channel when the program airs. Active promos are described in further detail below.
- A common problem is the baseball game overrun problem. VCRs and Digital Video Recorders (DVR) cut off the end of the baseball game
35 whenever the game runs over the advertised time slot. A TiVo tag is sent in the video stream indicating that the recording needs to continue. A TiVo tag is also sent telling the system to stop the recording because the game has ended.

- Boxing matches often end abruptly, causing VCRs and DVRs to record fill-in programs for the rest of the reserved time period. A TiVo tag is sent to indicate that the program has ended, telling the system to stop the recording.

5

- Referring to Fig. 19, advertisements are tagged so a locally or remotely stored advertisement might be shown instead of a national or out of the area advertisement. Within the video stream 1901, the program segment 1902 (commercial or other program segment) to be overlaid is tagged using techniques such as the TVTAG system described above. The TiVo tags tell the invention 1905 the start and end points of the old program segment 1902. A single tag 1903 can be added that tells the invention 1905 the duration of the old program segment 1902 or a tag is added at the beginning 1903 and end 1904 of the old program segment to indicate the start and end of the segment 1902. When the TiVo tag is detected, the invention 1905 finds the new program segment 1906 and simply plays it back in place of the old program segment 1902, reverting to the original program 1901 when playback is completed. The viewer 1907 never notices the transition.

10

15

20

There are three options at this point:

1) The system 1905 can continue to cache the original program, so if the viewer 1907 rewinds the program 1901 and plays it again, he sees the overlaid segment;

25

2) The old program segment 1902 is replaced in the cache too, so the viewer never sees the overlaid segment; or

30

3) The system caches the original segment 1902 and reinterprets the tags on playback. However, without intelligent tag prefetching, this only works correctly if the viewer backs up far enough so the system sees the first tag in the overlaid segment.

35

This problem is solved by adding the length of the old program segment to the start 1903 and end 1904 tag. Another approach is to match tags so that the start tag 1903 identifies the end tag 1904 to the system. The system 1905 knows that it should be looking for another tag when it fast forwards or rewinds over one of the tags. The pair of tags 1903, 1904 include a unique identifier. The system 1905 can then search ahead or

behind for the matching tag and replace the old program. There is a limit to the amount of time or length of frames that the system can conduct the prefetch. This can be included in the tag or standardized. Including the limit in the tag is the most flexible approach.

5

10 The program segment to be played back is selected based, for example, on locale, the time of day, program material, or on the preference engine (described in Application No. 09/422,121 owned by the Applicant). Using the preference engine, the appropriate program segment from local or server storage 1906 is selected according to the viewer's profile. The profile contains the viewer's viewing habits, program preferences, and other personal information. The stored program segments 1906 have program objects describing their features as well, which are searched for best match versus the preference vector.

15

20 Clearly, there must be a rotation mechanism among commercials to avoid ad burnout. The preference vector can be further biased by generating an error vector versus the program data for the currently viewed program, and using this error vector to bias the match against the commercial inventory on disk 1906. For example, if the viewer is watching a soap opera and the viewer's preference vector is oriented towards sports shows, then the invention will select the beer commercial in favor of the diaper commercial.

25 A tag can also be used to make conditional choices. The tag contains a preference weighting of its own. In this case, the preference weighting is compared to the preference vector and a high correlation causes the invention to leave the commercial alone. A low correlation invokes the method above.

30 NOTE: In all of these cases the system 1905 has more than enough time to make a decision. The structure of the pipeline routinely buffers 1/2 second of video, giving lots of time between input and output to change the stream. If more time is needed, add buffering to the pipeline. If playing back off disk, then the system creates the same time delay by reading ahead in the stream.

35 Also note that commercials can also be detected using the method described in Application No. 09/187,967 entitled "Analog Video Tagging and Encoding System," also owned by the Applicant. The same type of substitution described above can be used when tags described the aforementioned

application are used.

- With respect to Figs. 19 and 22, tags allow the incorporation of commercial "zapping." Since tags can be used to mark the beginning 1903 and ending 1904 points of a commercial, they can be skipped as well as preempted. The viewer simply presses the jump button 2205 on the remote control 2201. The system searches for the end tag and resumes playback at the frame following the frame associated with the tag. The number of commercials skipped is dependent upon the amount of video stream buffered.

Depending on the viewer's preset preferences, the system 1905 itself can skip commercials on live or prerecorded programs stored in memory 1906. Skipping commercials on live video just requires a larger amount of buffering in the pipeline as described above. Allowing the system to skip commercials on recorded programs presents the viewer with a continuous showing of the program without any commercial interruptions.

- Tags are added to program material to act as indexes. The viewer, for example, can jump to each index within the program by pressing the jump button 2205 on the remote control 2201.
- Tags are also used for system functions. As noted above, the system locally stores program material for its own use. The system 1905 must somehow receive the program material. This is done by tuning in to a particular channel at off hours. The system 1905 searches for the tag in the stream 1901 that tells it to start recording. The recording is comprised of a number of program segments delimited by tags 1903, 1904 that identify the content and possibly a preference vector. A tag at the end of the stream tells the system 1905 to stop recording. The program segments are stored locally 1906 and indexed for later use as described above.

The invention incorporates the following design points:

- The design provides for a clear separation of mechanism and policy.
- Internally, tags are viewed as abstract events which trigger policy modules. Mapping of received tag information to these internal abstractions is the

responsibility of the source pipeline object.

- Abstract tags are stored in the PesBuf stream as if they were just another segment. This allows the handling of arbitrary sized tags with precise timing information. It also allows tags to persist as part of recorded programs, so that proper actions are taken no matter when the program is viewed.
- Tags may update information about the current program, future programs, etc. This information is preserved for recorded programs.
- Tags can be logged as they pass through the system. It also possible to upload this information. It may not be necessary to preserve all information associated with a tag.
- Tags can be generated based on separate timelines. For example, using a network station log to generate tags based on time and network being viewed. Time-based tags are preserved in recorded streams.

Time-Based Tags

Referring to Fig. 20, time-based tags are handled by a Time-based Tag Recognizer 2012. This object 2012 listens for channel change events and, when a known network is switched to, attempts to retrieve a "time log" for that network. If one is present, the object 2012 builds a tag schedule based on the current time. As the time occurs for each tag, the object 2012 sends an event to the source object 2001 indicating the tag to be inserted. The source object 2001 inserts the tag into the next available position in the current PesBuf under construction. The next "available" position may be determined based on frame boundaries or other conditions.

The Role of the Source Object

The source object 2001 is responsible for inserting tags into the PesBuf stream it produces. This is assuming there are separate source objects for analog input and digital TV sources.

There are a number of different ways that tags may appear in an analog stream:

- Within the EDS field.
- Implicitly using the CC field.
- Modulated onto the VBI, perhaps using the ATVEF specification.
- Time Based

5

In a digital TV stream, or after conversion to MPEG from analog:

- In-band, using TiVo Tagging Technology.
- MPEG2 Private data channel.
- MPEG2 stream features (frame boundaries, etc.).
- Time-based tags.

10

15

The source object 2001 is not responsible for parsing the tags and taking any actions. Instead, the source object 2001 should solely be responsible for recognizing potential tags in the stream and adding them to the PesBuf stream.

Tag Recognition and Action

20

Conceptually, all tags may be broken up into two broad groups: those that require action upon reception, such as recording a program; and those that require action upon presentation, *i.e.*, when the program is viewed.

Reception Tag Handling

25

Tags that require action upon reception are handled as follows: a new Reception Tag Mechanism subclass 2003 of the TmkPushSwitch class 2002 is created. As input streams pass through this class 2003 between the source object 2001 and the program cache transform 2013, the class 2003 recognizes reception tags and takes appropriate actions.

30

Reception tags are generally handled once and then disabled.

Presentation Tag Handling

35

Tags that require actions upon presentation are handled as follows: a new Presentation Tag Mechanism subclass 2007 of the TmkPullSwitch class 2008 is created. As output streams pass through this class 2007 between the program cache transform 2013 and the sink object 2011, the class 2007 recognizes

presentation tags and takes appropriate actions.

Tag Policy Handling

- 5 Tag reception handling is only permitted if there is a TagReceptionPolicy object 2009 present for the current channel. Tag presentation handling is only permitted if there is a TagPresentationPolicy object 2010 for the source channel.

10 The TagPolicy objects describe which tags are to be recognized, and what actions are allowed.

15 When an input channel change occurs, the reception tag object is notified, and it fetches the TagReceptionPolicy object 2009 (if any) for that channel, and obeys the defined policy.

When an output channel change occurs, the presentation tag object is notified, and it fetches the TagPresentationPolicy object 2010 (if any) for that channel, and obeys the defined policy.

20 Tag Logging

25 The reception of tags may be logged into the database. This only occurs if a TagReceptionPolicy object 2009 is present, and the tag logging attribute is set. As an example, the logging attribute might be set, but no reception actions allowed to be performed. This allows passive logging of activity in the input stream.

Pipeline Processing Changes

30 It is important to support updates of information about the current showing. The following strategy is proposed:

- Whenever the input source is changed or a new showing starts, a copy is made of the showing object, and all further operations in the pipeline work off this copy.

35 - Update tags are reception tags; if permitted by policy, the copied showing object is updated.

- If the current showing is to be recorded, the copy of the showing object is saved with it, so that the saved program has the proper information saved with it.

- The original showing object is not modified by this process.

- The recorder must be cognizant of changes to the showing object, so that it doesn't, for instance, cut off the baseball game early.

10 Tag Interpretation vs. Tag State Machine

Tags are extremely flexible in that, once the TagPolicy object has been used to identify a valid tag, standardized abstract tags are interpreted by the Tag Interpreter 2005 and operational tags are executed by the TiVo Tag State Machine 2006. Interpreted tags trigger a predefined set of actions. Each set of actions have been preprogrammed into the system.

State machine tags are operational tags that do not carry executable code, but perform program steps. This allows the tag originator to combine these tags to perform customized actions on the TiVo system. State machine tags can be used to achieve the same results as an interpreted tag, but have the flexibility to dynamically change the set of actions performed.

25 Abstract Interpreted Tags

The set of available abstract tags is defined in a table called the Tag/Action table. This table is typically stored in a database object. There are a small number of abstract actions defined. These actions fall into three general categories:

- Viewer visible actions (may include interaction).
- Meta-information about the stream (channel, time, duration, etc.).
- TiVo control tags.

Tags which cause a change to the on-disk database, or cause implicit recording, must be validated. This is accomplished through control tags.

Viewer Visible Tags

- Menu

This tag indicates that the viewer is to be presented with a choice. The data associated with the tag indicates what the choice is, and other interesting data, such as presentation style. A menu has an associated inactivity timeout.

The idea of the menu tag is that the viewer is offered a choice. If the viewer isn't present, or is uninterested, the menu should disappear quickly. The menu policy may or may not be to pause the current program. The presentation of the menu does not have to be a list.

- Push Alternate Program Conditional

This tag indicates that some alternate program should be played if some condition is true. The condition is analyzed by the policy module. It may always be true.

- Pop Alternate Program Conditional

This tag reverts to the previous program. If a program ends, then the alternate program stack is popped automatically. All alternate programs are popped if the channel is changed or the viewer enters the TiVo Central menu area.

Alternate programs are a way of inserting arbitrary sequences into the viewed stream. The conditional data is not evaluated at the top level. Instead, the policy module must examine this data to make choices. This, for example, can be used to create "telescoping" ads.

- Show Indicator Conditional

This tag causes an indication to be drawn on the screen. Indicators are named, and the set of active indicators may be queried at any time. The tag or tag policy may indicate a timeout value at which time the indicator is derived.

- Clear Indicator Conditional

This tag causes an active indication to be removed. All indicators are cleared if the channel is changed or the viewer enters the TiVo Central menu area.

Indicators are another way to offer a choice to the viewer without interrupting program flow. They may also be used to indicate conditions in the stream that may be of interest. For example, "Active Promo" is created by providing a program object ID as part of the tag data, allowing that program to be selected. If the viewer hits a particular key while the indicator is up, then the program is scheduled for recording.

Meta-Information Tags

- Current Showing Information

This tag is a general bucket for information about the current showing. Each tag typically communicates one piece of information, such as the start time, end time, duration, etc. This tag can be used to "lengthen" a recording of an event.

- Future Showing Information

This tag is similar to the above, but contains information about a future showing. There are two circumstances of interest:

- The information refers to some showing already resident in the database. The database object is updated as appropriate.
- The information refers to a non-existent showing. A new showing object is created and initialized from the tag.

TiVo Control Tags

- Authorize Modification

This tag is generally encrypted with the current month's security key. The lifetime of the authorization is set by policy, probably to an hour or two. Thus, the tag needs to be continually rebroadcast if modifications to local TiVo system states are permitted.

The idea of this tag is to avoid malicious (or accidental) attacks using inherently insecure tag mechanisms such as EDS. If a network provides EDS information,

we first want to ensure that their tags are accurate and that attacks on the tag delivery system are unlikely. Then, we would work with that network to provide an authorization system that carouseled authorization tags on just that network. Unauthorized tags should never be inserted into the PES stream by the source object.

- Record Current Conditional

This tag causes the current program to be saved to disk starting from this point. The recording will cease when the current program ends.

- Stop Recording Current Conditional

This tag ceases recording of the current program.

- Record Future Conditional

A showing object ID is provided (perhaps just sent down in a Future Showing tag). The program is scheduled for recording at a background priority lower than explicit viewer selections.

- Cancel Record Future Conditional

A showing object ID is provided. If a recording was scheduled by a previous tag for that object, then the recording is canceled.

These tags, and the Future Showing tag, may be inserted in an encrypted, secure format. The source object will only insert these tags in the PES stream if they are properly validated.

One of the purposes of these tags is to automatically trigger recording of TiVo inventory, such as loopsets, advertisements, interstitials, etc. A later download would cause this inventory to be "installed" and available.

- Save File Conditional

This tag is used to pass data through the stream to be stored to disk. For instance, broadcast Web pages would be passed through

this mechanism.

- Save Object Conditional

- 5 This tag is used to pass an object through the stream to be stored to disk. Storing the object follows standard object updating rules.

10 The following is an example of an implementation using presentation tags inserted into the Closed Captioning (CC) part of a stream. The CC part of the stream was chosen because it is preserved when a signal is transmitted and digitized and decoded before it reaches the user's receiver. There are no guarantees on the rest of the VBI signal. Many of the satellite systems strip out everything except the closed captioning when encoding into MPEG-2.

- 15 There is a severe bandwidth limitation on the CC stream. The data rate for the CC stream is two 7-bit bytes every video frame. Furthermore, to avoid collision with the control codes, the data must start at 0x20, thus effectively limiting it to about 6.5-bit bytes (truncate to 6-bit bytes for simplicity). Therefore, the bandwidth is roughly 360 bits/second. This rate gets further reduced if the
- 20 channel is shared with real CC data. In addition, extra control codes need to be sent down to prevent CC-enabled televisions from attempting to display the TiVo tags as CC text.

Basic Tag Layout

- 25 This section describes how the tags are laid out in the closed captioning stream. It assumes a general familiarity with the closed captioning specification, though this is not crucial.

30 Making Tags Invisible

- 35 A TiVo Tag placed in a stream should not affect the display on a closed captioning enabled television. This is achieved by first sending down a "resume caption loading" command (twice for fault tolerance), followed by a string of characters that describes the tag followed by an "erase nondisplayed memory" command (twice for fault tolerance). What this does is to load text into offscreen memory, and then clear the memory. A regular TV with closed captioning enabled will not display this text (as per EIA-701 standard).

This works as long as the closed captioning decoder is not in "roll-up" or "scrolling" mode. In this mode, a "resume caption loading" command would cause the text to be erased. To solve this problem, TiVo Tags will be accepted and recognized even if they are sent to the second closed captioning channel. This way, even if closed captioning channel 1 is set up with scrolling text, we can still send the tag through closed captioning channel 2.

Tag Encoding

The text sent with a TiVo Tag consists of the letters "Tt", followed by a single character indicating the length of the tag, followed by the tag contents, followed by a CRC for the tag contents. The letters "Tt" are sufficiently unique that it is unlikely to encounter these in normal CC data. Furthermore, normal CC data always starts with a position control code to indicate where on the screen the text is displayed. Since we are not displaying onscreen, there is no need for this positioning data. Therefore, the likelihood of encountering a "Tt" immediately after a "resume caption loading" control code is sufficiently rare that we can almost guarantee that this combination is a TiVo tag (though the implementation still will not count on this to be true).

The single character indicating the length of the tag is computed by adding the tag length to 0x20. If the length is 3 characters, for example, then the length character used is 0x23 ('#'). So as not to limit the implementation to a length of 95 (since there are only 96 characters in the character set), the maximum length is defined as 63. If longer tags are needed, then an interpretation for the other 32 possible values for the length character can be added.

The possible values for the tag itself are defined in the Tag Types section below.

The CRC is the 16 bit CRC-CCITT (i.e., polynomial = $x^{16} + x^{12} + x^5 + 1$). It is placed in the stream as three separate characters. The first character is computed by adding 0x20 to the most significant six bits of the CRC. The next character is computed by adding 0x20 to the next six bits of the CRC. The last character is computed by adding 0x20 to the last four bits of the CRC.

Tag Types

This section details an example of a TiVo Tag. Note that every tag sequence begins with at least one byte indicating the tag type.

5 iPreview Tag

10 With respect to Fig. 17, an iPreview tag contains four pieces of information. The first is the 32 bit program ID of the program being previewed. The second contains how much longer the promotion is going to last. The third piece is where on the screen 1701 to place an iPreview alert 1702 and the last piece is what size iPreview alert to use.

15 The screen location for the iPreview alert is a fraction of the screen resolution in width and height. The X coordinate uses 9 bits to divide the width, so the final coordinate is given as: $X = (x_resolution/511) * xval$. If the xval is given as 10, on a 720 x 486 screen (using CCIR656 resolution), the X coordinate would be 14. The Y coordinate uses 8 bits to divide the height, so the final coordinate is given as: $Y = (y_resolution/255) * yval$. The X,Y coordinates indicate the location of the upper-left corner of the bug graphic.

20 If the value of X and Y are set to the maximum possible values (i.e., $x=511$, $y=255$), then this indicates that the author is giving the system the job of determining its position. The system will place the bug at a predetermined default position. The rationale for using the max values to indicate the default position is that it is never expected that a "real" position will be set to these values since that would put the entire bug graphic offscreen.

25 The size field is a four bit number that indicates what size any alert graphic should be. The 16 possible values of this field correspond to predefined graphic sizes that the settop boxes should be prepared to provide.

30 The timeout is a ten bit number indicating the number of frames left in the promotion. This puts a 34 second lifetime limit on this tag. If a promotion is longer, then the tag needs to be repeated. Note that the timeout was "artificially limited" to 10 bits to limit exposure to errors. This is to limit the effect it will have on subsequent commercials if an author puts a malformed timeout in the tag.

The version is a versioning number used to identify the promo itself. Instead of

bit-packing this number (and thus limiting it to 6 bits), the full closed captioning character set is used, which results in 96 possibilities instead of 64 (2^6). The version number thus needs to be within the range 0-95.

- 5 The reserved character is currently unused. This character needs to exist so that the control codes end up properly aligned on the 2-byte boundaries.

The first character of an iPreview tag is always "i".

- 10 All of the data fields are packed together on a bit boundary, and then broken into six bit values which are converted into characters (by adding 0x20) and transmitted. The order of the fields are as follows:

- 32 bits: program ID
- 15 • 9 bits: X location
- 8 bits: Y location
- 4 bits: graphic size
- 10 bits: timeout
- 1 character: version
- 20 • 1 character: reserved

The data fields total 66 bits which requires 11 characters to send + 1 character for version and 1 character for reserved. The exact contents of each character are:

- 25 1) 0x20 + ID[31:26]
- 2) 0x20 + ID[25:20]
- 3) 0x20 + ID[19:14]
- 4) 0x20 + ID[13:8]
- 5) 0x20 + ID[7:2]
- 30 6) 0x20 + ID[1:0] X[8:5]
- 7) 0x20 + X[4:0] Y[7]
- 8) 0x20 + Y[6:1]
- 9) 0x20 + Y[0] size[3:0]
- 10) 0x20 + Y[0] size[3:0] timeout[9]
- 35 11) 0x20 + timeout[8:3]
- 12) 0x20 + timeout[2:0]
- 13) 0x20 + version
- 14) reserved

Including the first character "i", the length of the iPreview tag is 14 characters + 3 CRC characters. With the tag header (3 characters), this makes a total length of 20 characters which can be sent down over 10 frames. Adding another 4 frames
 5 for sending "resume caption loading" twice and "erase nondisplayed memory" twice means an iPreview tag will take 14 frames (0.47 seconds) to broadcast.

A complete iPreview tag consists of:

10 Resume caption loading Resume caption loading T t 1 (0x20 + 17 = 0x31 = 0110001= "1") i <13 character iPreview tag> 3 character CRC Erase
 nondisplayed memory Erase nondisplayed memory

Parity debugging character

15 Currently, the parity bit is being used as a parity bit. However, since a CRC is already included, there is no need for the error-checking capabilities of the parity bit. Taking this a step further, the parity bit can be used in a clever way. Since a closed captioning receiver should ignore any characters with an incorrect parity bit,
 20 a better use of the limited bandwidth CC channel can be had by intentionally using the wrong parity. This allows the elimination of the resume caption loading and erase nondisplayed memory characters, as well as making it easier to "intersperse" TiVo tags among existing CC data.

iPreview Viewer Interaction

25 Referring to Figs. 17, 20, 21 and 22, the iPreview tag causes the Tag Interpreter 2005 to display the iPreview alert 1702 on the screen 1701. The iPreview alert 1702 tells the viewer that an active promo is available and the viewer can tell the TiVo system to record the future showing. The viewer reacts to the iPreview alert
 30 1702 by pressing the select button 2204 on the remote control 2201.

The Tag Interpreter 2005 waits for the user input. Depending on the viewer's preset preferences, the press of the select button 2204 results in the program automatically scheduled by the Tag Interpreter 2005 for recording, resulting in a
 35 one-touch record, or the viewer is presented with a record options screen 2101. The viewer highlights the record menu item 2102 and presses the select button 2204 to have the program scheduled for recording.

The tag itself has been interpreted by the Tag Interpreter 2005. The Tag Interpreter 2005 waits for any viewer input through the remote control 2201. Once the viewer presses the select button 2204, the Tag Interpreter 2005 tells the TiVo system to schedule a recording of the program described by the 32 bit program ID in the iPreview tag.

With respect to Figs. 20, 22, and 23, the iPreview tag is also used for other purposes. Each use is dictated by the context of the program material and the screen icon displayed. Obviously the system cannot interpret the program material, but the icon combined with the program ID tell the Tag Interpreter 2005 what action to take. Two examples are the generation of a lead and a sale.

The process of generating a lead occurs when, for example, a car ad is being played. An iPreview icon appears 2301 on the screen and the viewer knows that he can press the select button 2204 to enter an interactive menu.

A menu screen 2302 is displayed by the Tag Interpreter 2005 giving the user the choice to get more information 2303 or see a video of the car 2304. The viewer can always exit by pressing the live TV button 2202. If the viewer selects get more information 2303 with the up and down arrow button 2203 select button 2204, then the viewer's information is sent to the manufacturer 2305 by the Tag Interpreter 2005, thereby generating a lead. The viewer returns to the program by pressing the select button 2204.

Generating a sale occurs when a product, *e.g.*, a music album ad, is advertised. The iPreview icon 2301 appears on the screen. The viewer presses the select button 2204 and a menu screen 2307 is displayed by the Tag Interpreter 2005.

The menu screen 2307 gives the viewer the choice to buy the product 2308 or to exit 2309. If the viewer selects yes 2308 to buy the product, then the Tag Interpreter 2005 sends the order to the manufacturer with the viewer's purchase information 2310. If this were a music album ad, the viewer may also be presented with a selection to view a music video by the artist.

Whenever the system returns the viewer back to the program, it returns to the exact point that the viewer had originally exited from. This gives the viewer a sense of continuity.

The concept of redirection is easily expanded to the Internet. The iPreview icon will appear as described above. When the viewer presses the select button 2204 on the remote control 2201, a Web page is then displayed to the viewer. The viewer then interacts with the Web page and when done, the system returns the viewer back to the program that he was watching at the exact point from which the viewer had exited.

Using the preference engine as noted above, the information shown to the viewer during a lead or sale generation is easily geared toward the specific viewer. The viewer's viewing habits, program preferences, and personal information are used to select the menus, choices, and screens presented to the viewer. Each menu, choice and screen has an associated program object that is compared to the viewer's preference vector.

For example, if a viewer is male and the promo is for Chevrolet, then when the viewer presses the select button, a still of a truck is displayed. If the viewer were female, then a still of a convertible would be displayed.

Note that the Tag State Machine 2006 described below is fully capable of performing the same steps as the Tag Interpreter 2005 in the above examples.

The TiVo Tag State Machine

Referring again to Fig. 20, a preferred embodiment of the invention provides a Tag State Machine (TSM) 2006 which is a mechanism for processing abstract TiVo tags that may result in viewer-visible actions by the TiVo Receiver.

A simple example is the creation of an active promo. As demonstrated above, an active promo is where a promotion for an upcoming show is broadcast and the viewer is immediately given the option of having the TiVo system record that program when it actually is broadcast.

Hidden complexities underlie this simple example: some indicator must be generated to alert the viewer to the opportunity; the indicator must be brought into view or removed with precision; accurate identification of the program in question must be provided; and the program within which the active promo appears may be viewed at a very different time than when it was broadcast.

Creation and management of the TiVo tags is also challenging. It is important to cause as little change as possible to existing broadcast practices and techniques. This means keeping the mechanism as simple as possible for both ease of integration into the broadcast stream and for robust and reliable operation.

5

Principles of Tags

10

As previously noted, it is assumed that the bandwidth available for sending tags is constrained. For example, the VBI has limited space available which is under heavy competition. Even in digital television signals, the amount of out-of-band data sent will be small since most consumers of the signal will be mainly focused on television programming options.

15

A tag is then a simple object of only a few bytes in size. More complex actions are built by sending multiple tags in sequence.

20

The nature of broadcast delivery implies that tags will get lost due to signal problems, sunspots, etc. The TSM incorporates a mechanism for handling lost tags, and insuring that no unexpected actions are taken due to lost tags.

25

In general, viewer-visible tag actions are relevant only to the channel on which they are received; it is assumed that tag state is discarded after a channel change.

30

Physical tags are translated into abstract tags by the source object 1901 receiving the physical tag. Tags are not "active agents" in that they carry no executable code; functioning the TSM may result in viewer-visible artifacts and changes, but the basic operation of the TiVo receiver system will remain unaffected by the sequence of tags. If tags could contain executable code, such as the Java byte streams contemplated by the ATVEF, the integrity of the TiVo viewing experience might be compromised by poorly written or malicious software.

35

All tag actions are governed by a matching policy object matched to the current channel. Any or all actions may be enabled or disabled by this object; the absence of a policy object suppresses all tag actions.

The Basic Abstract Tag

All abstract tags have a common infrastructure. The following components are present in any abstract tag:

- Tag Type (1 byte)

The type 0 is disallowed. The type 255 indicates an "extension" tag, should more than 254 tag values be required at some future time.

- Tag Sequence (1 byte)

This unsigned field is incremented for each tag that is part of a sequence. Tags which are not part of a sequence must have this field set to zero. A tag sequence of one indicates the start of a new sequence; a sequence may be any length conceptually, but it will be composed of segments of no more than 255 tags in order.

Each tag type has an implicit sequence length (which may be zero); the sequence number is introduced to handle dropouts or other forms of tag loss in the stream. In general, if a sequence error occurs, the entire tag sequence is discarded and the state machine reset.

Tags should be checksummed in the physical domain. If the checksum doesn't match, the tag is discarded by the source object. This will result in a sequence error and reset of the state machine.

- Tag Timestamp (8 bytes)

This is the synchronous time within the TV stream at which the tag was recognized. This time is synchronous to all other presentation times generated by the TiVo Receiver. This component is never sent, but is generated by the receiver itself.

- Tag Data Length (2 bytes)

This is the length of any data associated with the tag. The interpretation of this data is based on the tag type. The physical domain translator should perform some minimal error checking on the data.

The Tag State Machine

The TSM is part of the Tag Presentation Mechanism, which is in-line with video playback.

5

Conceptually, the TSM manages an abstract stack of integer values with at least 32 bits of precision, or sufficient size to hold an object iD. The object iD is abstract, and may or may not indicate a real object on the TiVo Receiver - it may otherwise need to be mapped to the correct object. The stack is limited in size to

10

255 entries to limit denial-of-service attacks.

The TSM also manages a pool of variables. Variables are named with a 2-byte integer. The variable name 0 is reserved. "User" variables may be manipulated by tag sequences; such variables lie between 1 and $2^{15}-1$. "System" variables are maintained by the TSM, and contain values about the current TiVo Receiver, such as: the current program object ID; the TSM revision; and other useful information. These variables have names between 2^{15} and $2^{16}-1$. The number of user variables may be limited within a TSM; a TSM variable indicates what this limit is.

15

20

The tag data is a sequence of TSM commands. Execution of these commands begins when the tag is recognized and allowed. TSM commands are byte oriented and certain commands may have additional bytes to support their function.

25

The available TSM commands may be broken down into several classes:

Data Movement Commands

30

push_byte - push the byte following the command onto the stack.
push_short - push the short following the command onto the stack.
push_word - push the word following the command onto the stack.

Variable Access Commands

35

push_var - push the variable named in the 16-bit quantity following the command.
pop_var - pop into the variable named in the 16-bit quantity following the

command.

copy_var - copy into the variable named in the 16-bit quantity following the command from the stack.

5 Stack Manipulation Commands

swap - swap the top two stack values.

pop - toss the top stack value.

10 Arithmetic Commands

add_byte - add the signed byte following the command to the top of stack.

add_short - add the signed short following the command to the top of stack.

add_word - add the signed word following the command to the top of stack.

15 and - and the top and next stack entries together, pop the stack and push the new value.

or - or the top and next stack entries together, pop the stack and push the new value.

20 Conditional Commands

(Unsigned comparisons only)

25 brif_zero - branch to the signed 16-bit offset following the command if the top of stack is zero.

brif_nz - branch to the signed 16-bit offset following the command if the top of stack is not zero.

brif_gt - branch to the signed 16-bit offset following the command if the top of stack is greater than the next stack entry.

30 brif_ge - branch to the signed 16-bit offset following the command if the top of stack is greater than or equal to the next stack entry.

brif_le - branch to the signed 16-bit offset following the command if the top of stack is less than or equal to the next stack entry.

35 brif_lt - branch to the signed 16-bit offset following the command if the top of stack is less than the next stack entry.

brif_set - branch to the signed 16-bit offset following the command if there are bits set when the top and next stack entries are ANDed together.

Action Commands

exec - execute tag action on the object ID named on top of stack.

fin - terminate tag taking no action.

5

System Variables

32768 (TAG) - value of current tag.

⋮

10 Times in GMT:

32769 (YEAR) - current year (since 0).

32770 (MONTH) - current month (1-12).

32771 (DAY) - day of month (1-31).

15 32772 (WDAY) - day of week (1-7, starts Sunday).

32773 (HOUR) - hour of the day (0-23).

32774 (MIN) - minute of the hour (0-59).

32775 (SEC) - seconds of the minute (0-59).

20 TiVo Receiver State:

32800 (SWREL) - software release (in x.x.x notation in bytes).

32801 (NTWRK) - object ID of currently tuned network.

32802 (PRGRM) - object ID of currently tuned program.

25 32803 (PSTATE) - current state of output pipeline:

0 - normal playback

1 - paused

2 - slo-mo

10 - rewind speed 1

30 11 - rewind speed 2

...

20 - ff speed 1

21 - ff speed 2

...

35

Tag Execution State:

32900 (IND) - indicator number to display or take down.

32901 (PDURING) - state of the pipeline while tag is executing.

32902 (ALTP) - alternate program object ID to push on play stack.

32903 (SELOBJ) - program object ID to record if indicator selected.

5 33000 (MENU1) - string object number for menu item 1.

33001 (MENU2) - string object number for menu item 2.

...

33009 (MENU10) - string object number for menu item 9.

10 33100 (PICT1) - picture object number for menu item 1.

33101 (PICT2) - picture object number for menu item 2.

...

33109 (PICT10) - picture object number for menu item 10.

15 33200 (MSELOBJ1) - program object ID to record if menu item selected.

33201 (MSELOBJ2) - program object ID to record if menu item selected.

...

33209 (MSELOBJ10) - program object ID to record if menu item selected.

20 Tags

- Push Alternate Program

- Pop Alternate Program (auto-pop at end of program)

- Raise Indicator

25 - Lower Indicator

- Menu

Tag Execution Policy

30 Execution policy is determined by the TSM. Some suggestions are:

- Menus

35 Menus are laid out as per standard TiVo menu guidelines. In general, menus appear over live video. Selection of an item typically invokes the record dialog. It may be best to pause the pipeline during the menu operation.

- Indicators

With respect to Figs. 17 and 22, indicators 1702 are lined up at the bottom of the display as small icons. During the normal viewing state, the up arrow and down arrow keys 2203 on the remote control 2201 do nothing. For indicators, up arrow 2203 circles through the indicators to the left, down arrow to the right. The selected indicator has a small square drawn around it. Pushing select 2204 initiates the action. New indicators are by default selected; if an indicator is removed, the previously selected indicator is highlighted, if any.

- Alternate Programs

Alternate programs should appear as part of the video stream, and have full ff/rew controls. The skip to live button 2202 pops the alternate program stack to empty first.

One skilled in the art will readily appreciate that although the closed caption stream is specifically mentioned above, other transport methods can be used such as the EDS fields, VBI, MPEG2 private data channel, etc.

Although the invention is described herein with reference to the preferred embodiment, one skilled in the art will readily appreciate that other applications may be substituted for those set forth herein without departing from the spirit and scope of the present invention. Accordingly, the invention should only be limited by the Claims included below.

CLAIMS

1. A process for frame specific tagging of television audio and video
5 broadcast streams with tag translation at a receiver, comprising the steps of:

providing a storage device on said receiver;

inserting tags into said broadcast stream;

tuning said receiver to said broadcast stream;

receiving said broadcast stream at said receiver;

10 storing said broadcast stream on said storage device;

detecting said tags in said broadcast stream;

processing said tags;

displaying program material in said broadcast stream from said storage
device to a viewer;

15 wherein said processing step performs the appropriate actions in
response to said tags; and

wherein said tags include command and control information.

2. The process of claim 1, wherein tags indicate the start and end points of a
20 program segment.

3. The process of claim 2, wherein said displaying step skips over said
program segment in response to the viewer pressing a button on a remote input
device.

4. The process of claim 2, wherein said displaying step automatically skips
said program segment.

5. The process of claim 1, wherein said processing step displays a menu to
30 the viewer based on information included in a tag.

6. The process of claim 1, wherein said processing step records the current
program in said broadcast stream on said storage device based on information
included in a tag.

7. The process of claim 1, said processing step further comprising the steps
of:

displaying multiple icons to the viewer;

accepting viewer input information;
allowing the viewer to scroll through said multiple icons;
selecting a particular icon based on the viewer's input; and
performing an action associated with the selected icon.

5

8. The process of claim 1, further comprising the steps of:
wherein said processing step displays an icon to the viewer based on
information included in a tag;

10

accepting viewer input information;
interacting with the viewer based on the tag information;
wherein said displaying step saves the exit point in the program material;

and

wherein the viewer is returned to said exit point upon completion of any
interaction.

15

9. The process of claim 8, further comprising the steps of:
presenting a plurality of menus to the viewer for generating a lead; and
forwarding the viewer's contact information to a third party upon viewer
approval.

20

10. The process of claim 8, further comprising the steps of:
presenting a plurality of menus to the viewer for generating a sale of an
advertised product or service; and
forwarding the viewer's purchase information to the proper merchant.

25

11. The process of claim 8, further comprising the step of:
presenting a set of program recording options to the viewer; and
scheduling the viewer's recording preferences.

30

12. The process of claim 8, further comprising the step of:
presenting the content of a Web site's Web page to the viewer in
response to the viewer's input; and
wherein the viewer is allowed to interact with said Web site.

35

13. The process of claim 1, wherein said tags allow a system administrator to
remotely configure said receiver.

14. The process of claim 1, further comprising the steps of:

marking indexes in said program material based on tag information; and
jumping to an index selected by the viewer.

5 15. A process for scheduling the recording of a television program via an
advertisement in a television broadcast stream, comprising the steps of:

receiving said television broadcast stream;

playing a promotional advertisement in said television broadcast stream
for a future showing of a program;

10 displaying an icon notifying the viewer that said program is available to
record;

accepting the viewer's single key press from a remote input device; and
scheduling the recording of said program.

15 16. The process of claim 15, wherein said icon is displayed based on a tag
inserted into said television broadcast stream.

17. The process of claim 15, further comprising the step of:

providing a storage device; and

20 wherein said program is stored on said storage device when the
scheduled time arrives.

18. A process for the automatic replacement of program segments in a
multimedia television broadcast stream at a receiver, comprising the steps of:

receiving said multimedia television broadcast stream;

25 detecting the start and end points of an old program segment in said
broadcast stream;

providing a plurality of new program segments; and

substituting said old program segment with a new program segment
during playback of said broadcast stream to a viewer.

30 19. The process of claim 18, wherein said detecting step searches for tags
inserted into said broadcast stream denoting the start and end points of program
segments.

35 20. The process of claim 19, wherein said tags are located in the closed
caption area of said broadcast stream.

21. The process of claim 18, further comprising the step of:

providing a storage device on said receiver; and
wherein said new program segments are stored on said storage device.

5 22. The process of claim 21, further comprising the steps of:
 receiving new program segments via said broadcast stream; and
 storing said new program segments on said storage device.

10 23. The process of claim 18, wherein said new program segments are stored
 at a remotely accessible location.

15 24. The process of claim 18, wherein said new program segment to be
 played back is selected based on criteria such as: locale, the time of day,
 program material, the viewer's viewing habits, the viewer's program preferences,
 or the viewer's personal information.

20 25. The process of claim 24, wherein said criteria may result in the old program
 segment not being substituted.

25 26. The process of claim 24, wherein said new program segments have
 program objects describing their features which are used to select the best
 matching new program segment.

30 27. The process of claim 18, wherein a rotation mechanism is used when
 selecting said new program segments to avoid ad burnout.

35 28. An apparatus for frame specific tagging of television audio and video
 broadcast streams with tag translation at a receiver, comprising:

 a storage device on said receiver;

 a module for inserting tags into said broadcast stream;

 a module for tuning said receiver to said broadcast stream;

 a module for receiving said broadcast stream at said receiver;

 a module for storing said broadcast stream on said storage device;

 a module for detecting said tags in said broadcast stream;

 a module for processing said tags;

 a module for displaying program material in said broadcast stream from
said storage device to a viewer;

 wherein said processing module performs the appropriate actions in
response to said tags; and

wherein said tags include command and control information.

29. The apparatus of claim 28, wherein tags indicate the start and end points of a program segment.

5

30. The apparatus of claim 29, wherein said displaying module skips over said program segment in response to the viewer pressing a button on a remote input device.

10 31. The apparatus of claim 29, wherein said displaying module automatically skips said program segment.

32. The apparatus of claim 28, wherein said processing module displays a menu to the viewer based on information included in a tag.

15

33. The apparatus of claim 28, wherein said processing module records the current program in said broadcast stream on said storage device based on information included in a tag.

20 34. The apparatus of claim 28, said processing module further comprising:
a module for displaying multiple icons to the viewer;
a module for accepting viewer input information;
a module for allowing the viewer to scroll through said multiple icons;
a module for selecting a particular icon based on the viewer's input; and
25 a module for performing an action associated with the selected icon.

35. The apparatus of claim 28, further comprising:
wherein said processing module displays an icon to the viewer based on information included in a tag;

30 a module for accepting viewer input information;
a module for interacting with the viewer based on the tag information;
wherein said displaying module saves the exit point in the program material; and

35 wherein the viewer is returned to said exit point upon completion of any interaction.

36. The apparatus of claim 35, further comprising:
a module for presenting a plurality of menus to the viewer for generating a

lead; and

a module for forwarding the viewer's contact information to a third party upon viewer approval.

- 5 37. The apparatus of claim 35, further comprising:
 a module for presenting a plurality of menus to the viewer for generating a sale of an advertised product or service; and
 a module for forwarding the viewer's purchase information to the proper merchant.

- 10 38. The apparatus of claim 35, further comprising:
 a module for presenting a set of program recording options to the viewer;
 and
 a module for scheduling the viewer's recording preferences.

- 15 39. The apparatus of claim 35, further comprising:
 a module for presenting the content of a Web site's Web page to the viewer in response to the viewer's input; and
 wherein the viewer is allowed to interact with said Web site.

- 20 40. The apparatus of claim 28, wherein said tags allow a system administrator to remotely configure said receiver.

- 25 41. The apparatus of claim 28, further comprising:
 a module for marking indexes in said program material based on tag information; and
 a module for jumping to an index selected by the viewer.

- 30 42. An apparatus for scheduling the recording of a television program via an advertisement in a television broadcast stream, comprising:
 a module for receiving said television broadcast stream;
 a module for playing a promotional advertisement in said television broadcast stream for a future showing of a program;
 a module for displaying an icon notifying the viewer that said program is available to record;
35 a module for accepting the viewer's single key press from a remote input device; and
 a module for scheduling the recording of said program.

43. The apparatus of claim 42, wherein said icon is displayed based on a tag inserted into said television broadcast stream.

5 44. The apparatus of claim 42, further comprising:
a storage device; and
wherein said program is stored on said storage device when the scheduled time arrives.

10 45. A apparatus for the automatic replacement of program segments in a multimedia television broadcast stream at a receiver, comprising:

a module for receiving said multimedia television broadcast stream;
a module for detecting the start and end points of an old program segment in said broadcast stream;
15 a module for providing a plurality of new program segments; and
a module for substituting said old program segment with a new program segment during playback of said broadcast stream to a viewer.

20 46. The apparatus of claim 45, wherein said detecting module searches for tags inserted into said broadcast stream denoting the start and end points of program segments.

25 47. The apparatus of claim 46, wherein said tags are located in the closed caption area of said broadcast stream.

48. The apparatus of claim 45, further comprising:
a storage device on said receiver; and
wherein said new program segments are stored on said storage device.

30 49. The apparatus of claim 48, further comprising:
a module for receiving new program segments via said broadcast stream;
and
a module for storing said new program segments on said storage device.

35 50. The apparatus of claim 45, wherein said new program segments are stored at a remotely accessible location.

51. The apparatus of claim 45, wherein said new program segment to be

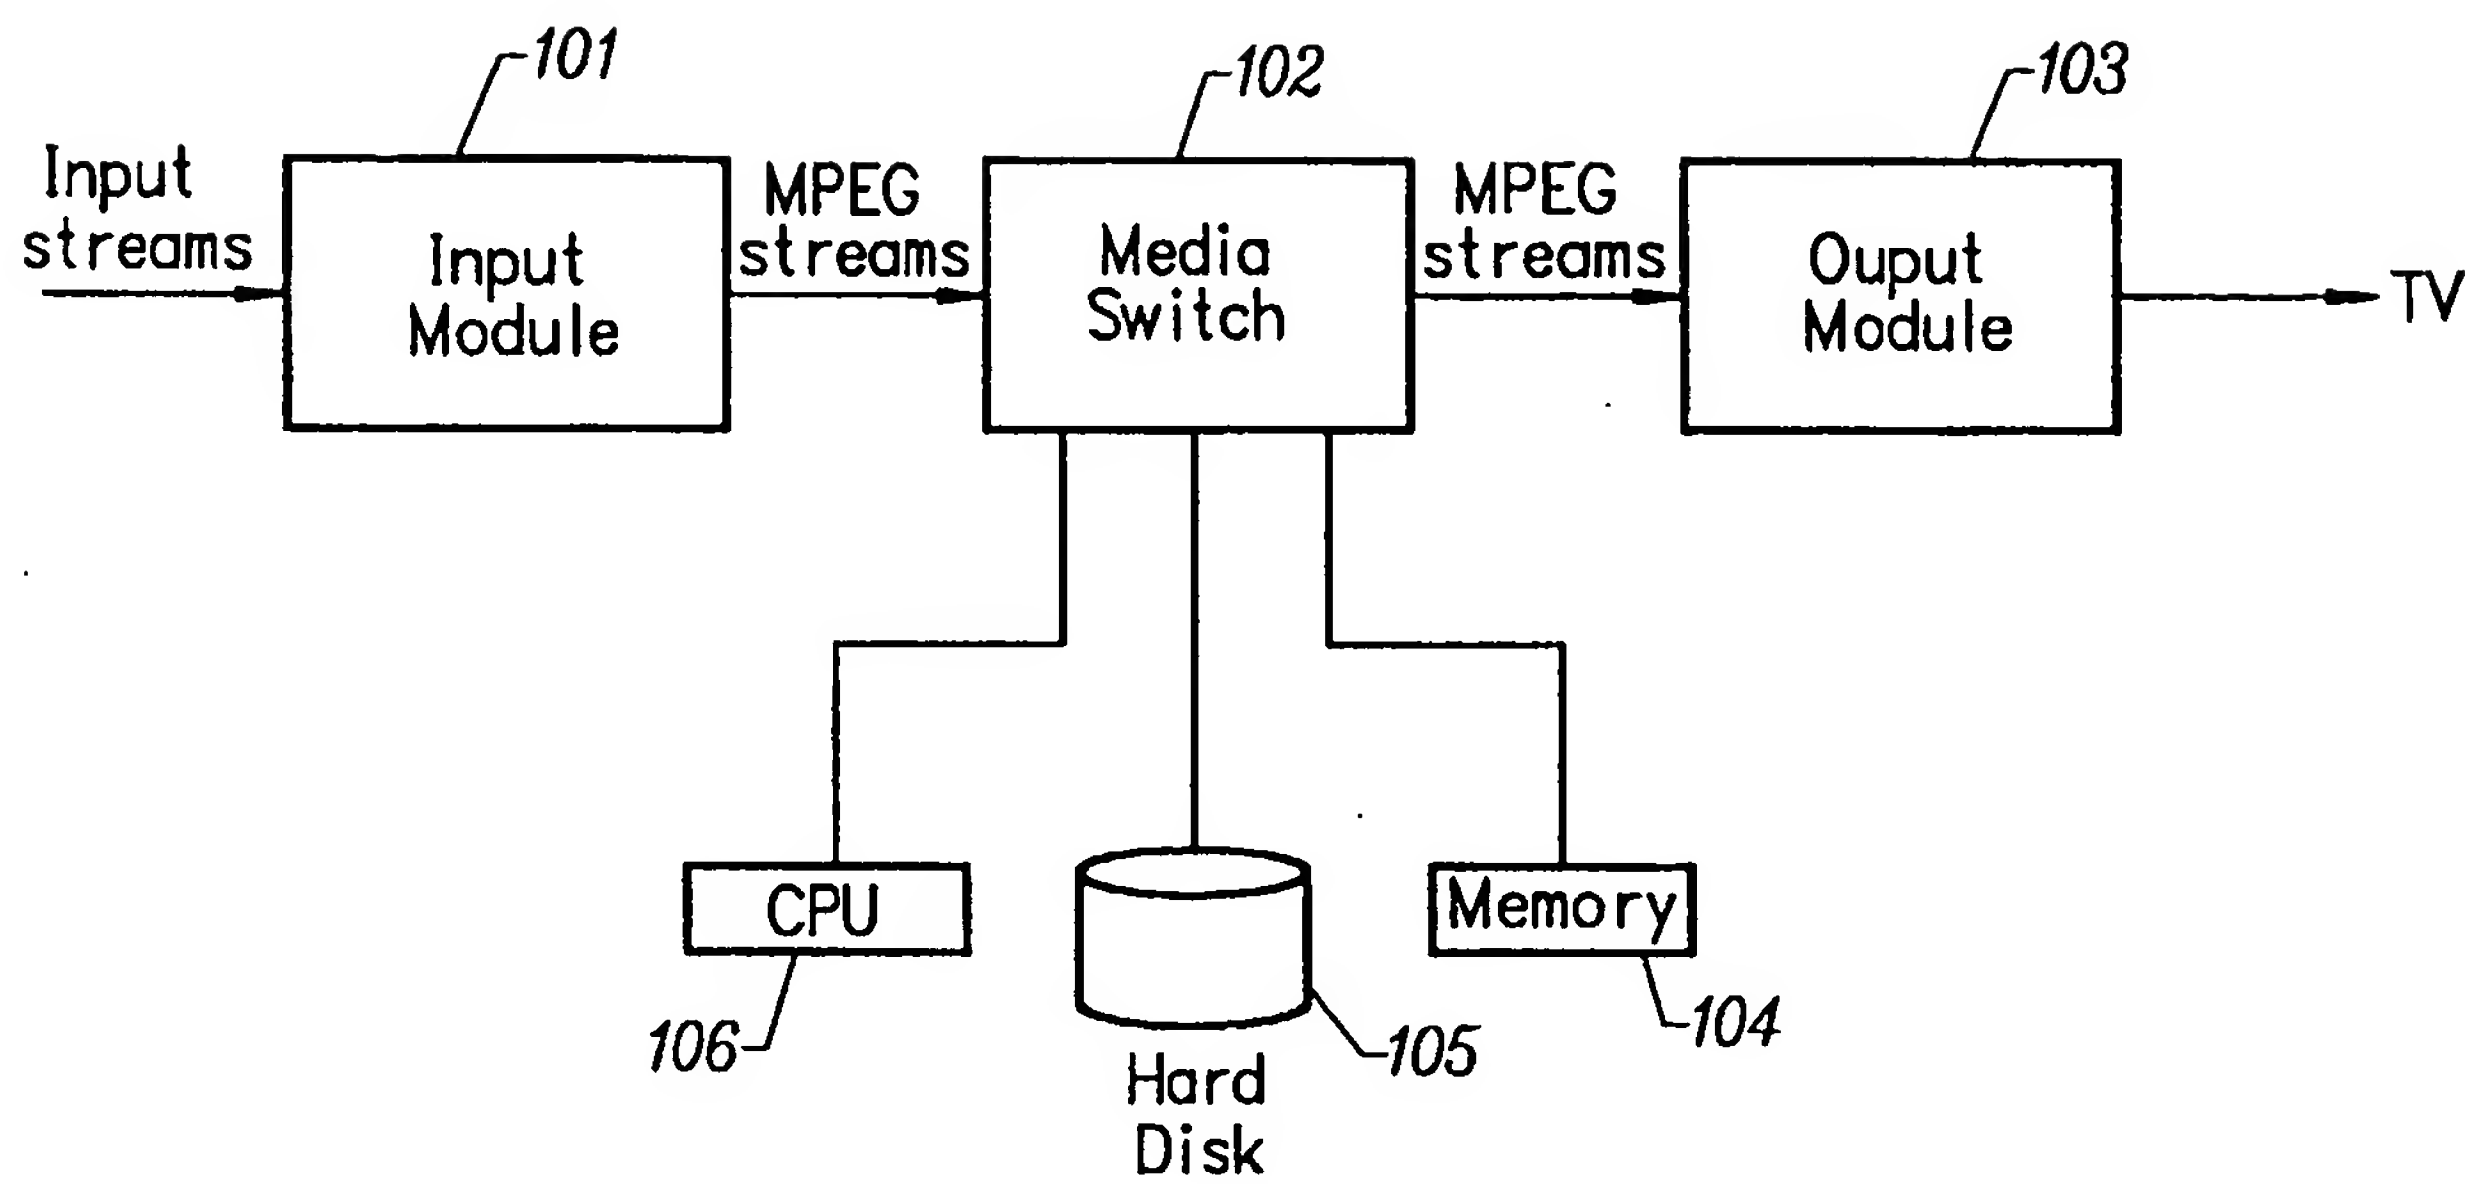
played back is selected based on criteria such as: locale, the time of day, program material, the viewer's viewing habits, the viewer's program preferences, or the viewer's personal information.

5 52. The apparatus of claim 51, wherein said criteria may result in the old program segment not being substituted.

10 53. The apparatus of claim 51, wherein said new program segments have program objects describing their features which are used to select the best matching new program segment.

54. The apparatus of claim 45, wherein a rotation mechanism is used when selecting said new program segments to avoid ad burnout.

1/22

*FIG. 1*

2/22




FIG. 2




FIG. 3

4/22




FIG. 4




FIG. 5

5/22




FIG. 6

6/22





FIG. 7

7/22

*FIG. 8*

8/22

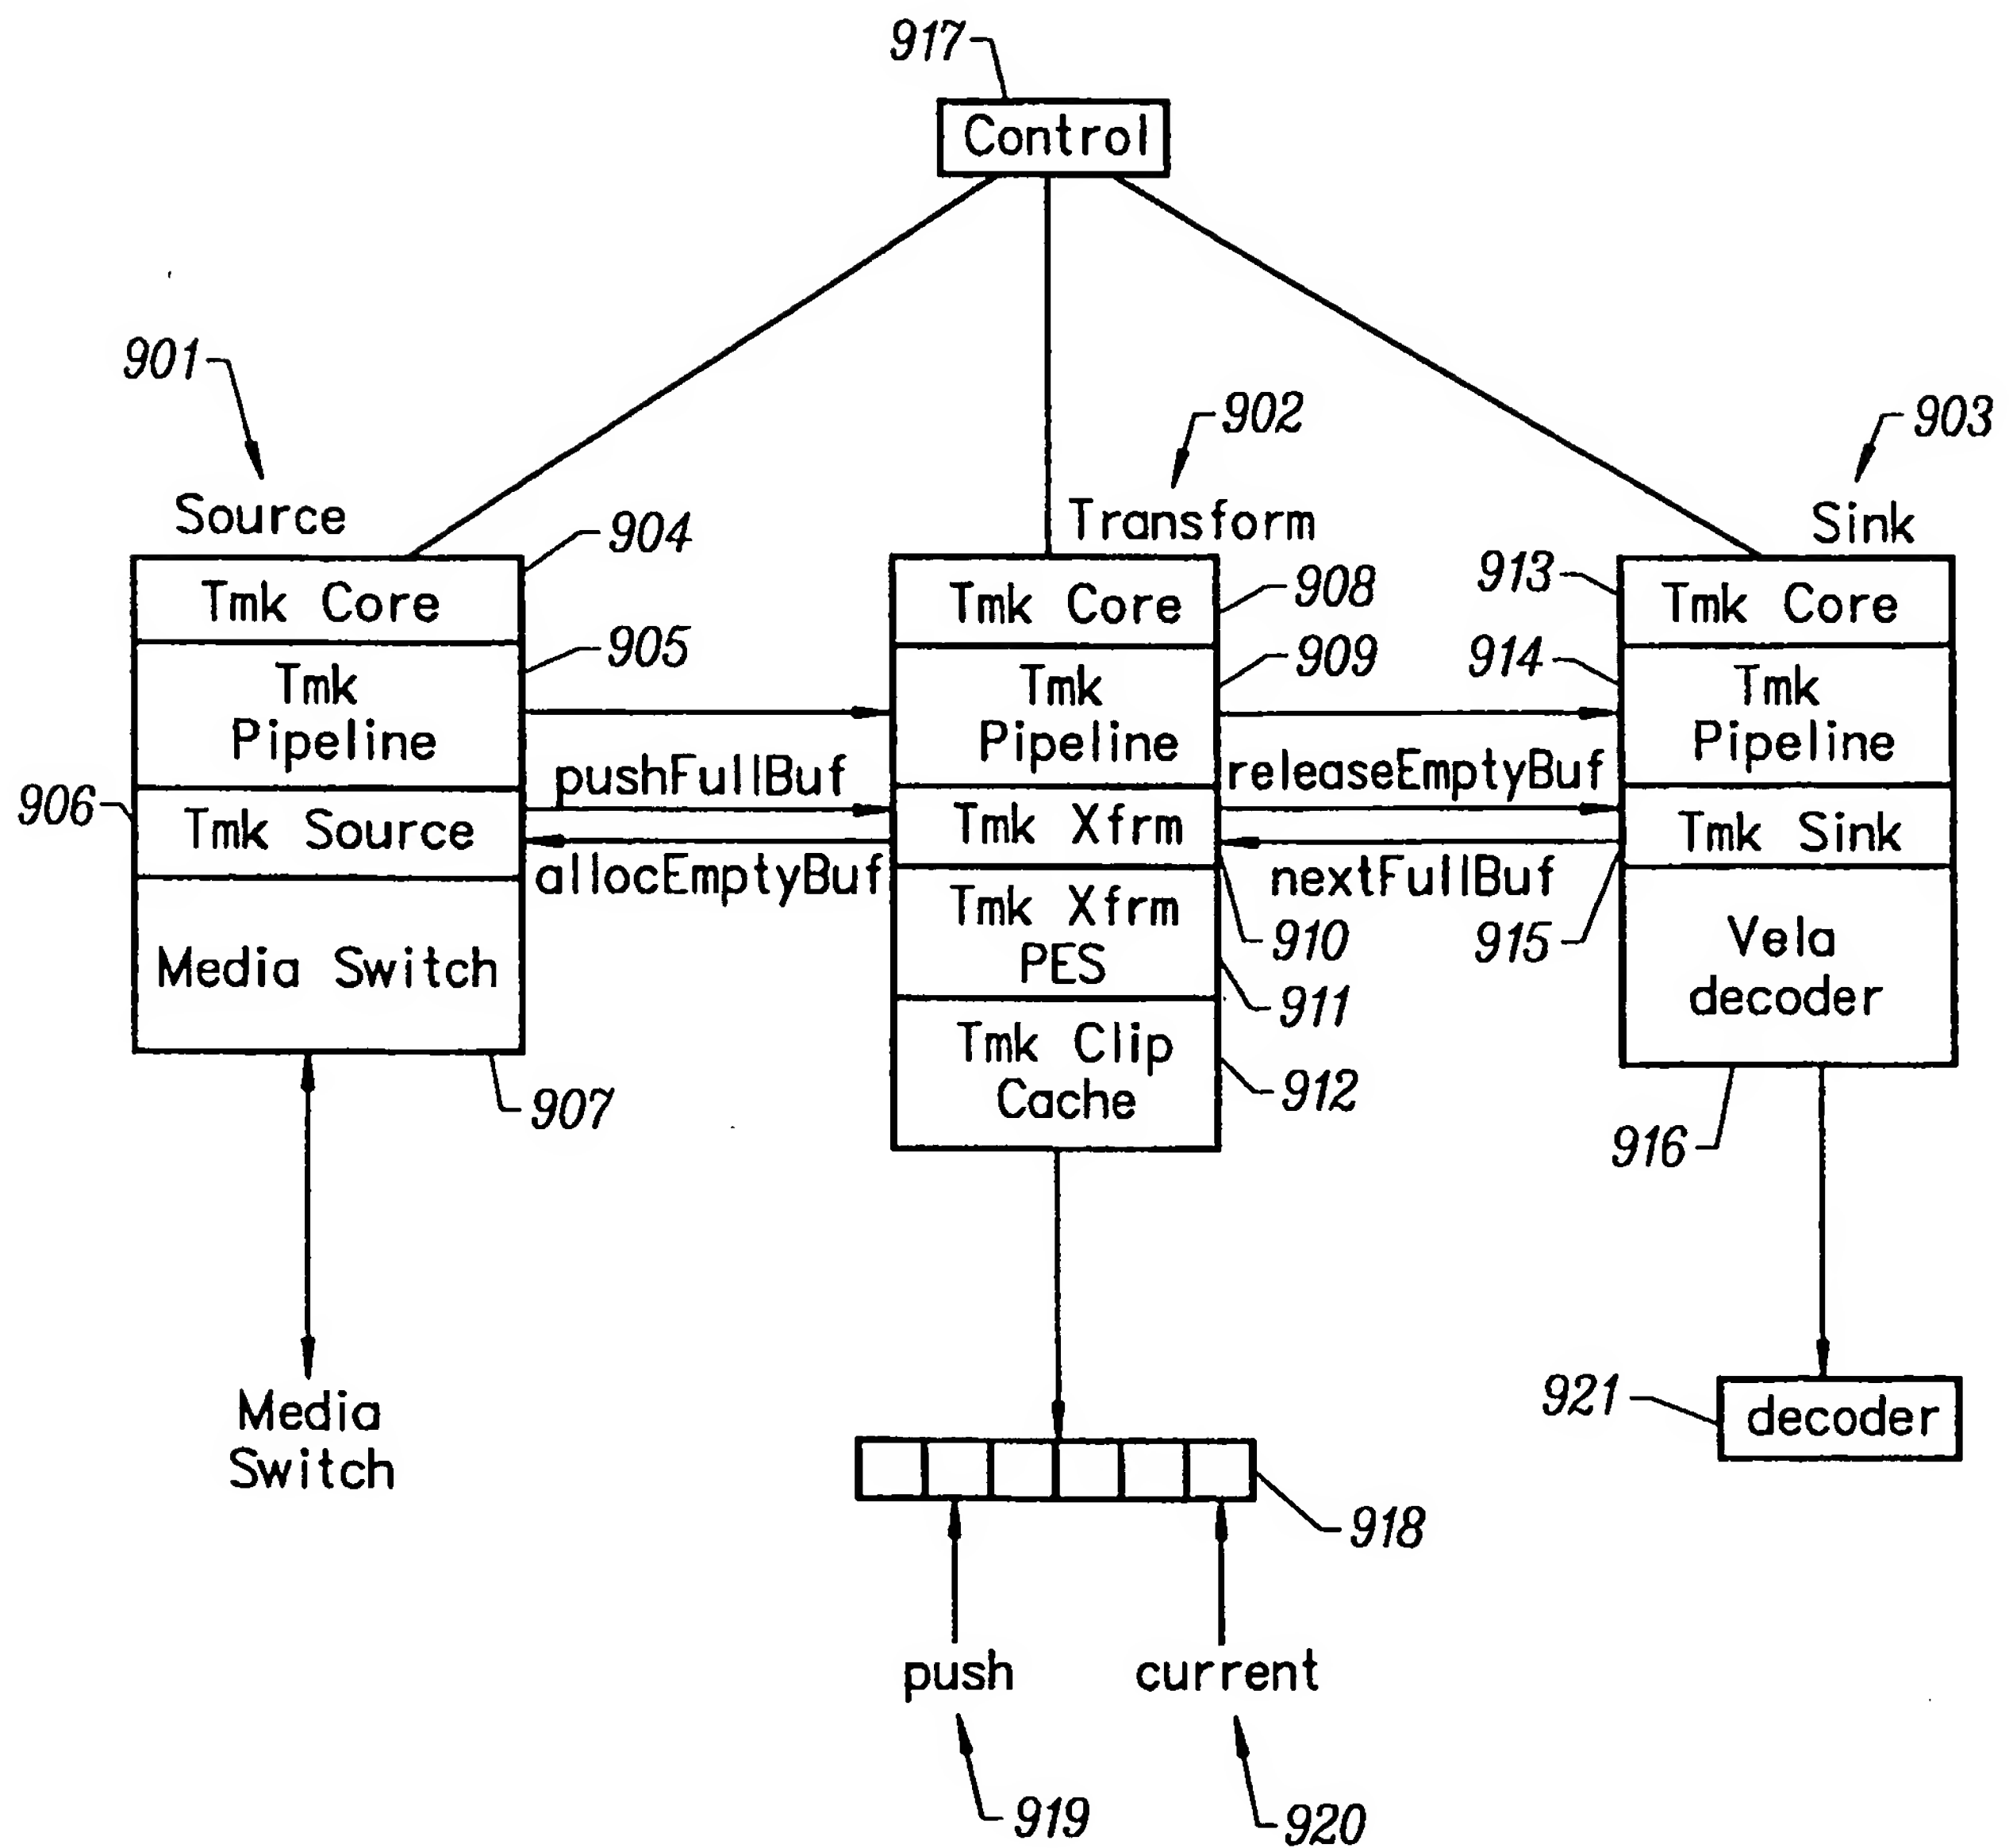



FIG. 9

9/22

*FIG. 10*

10/22

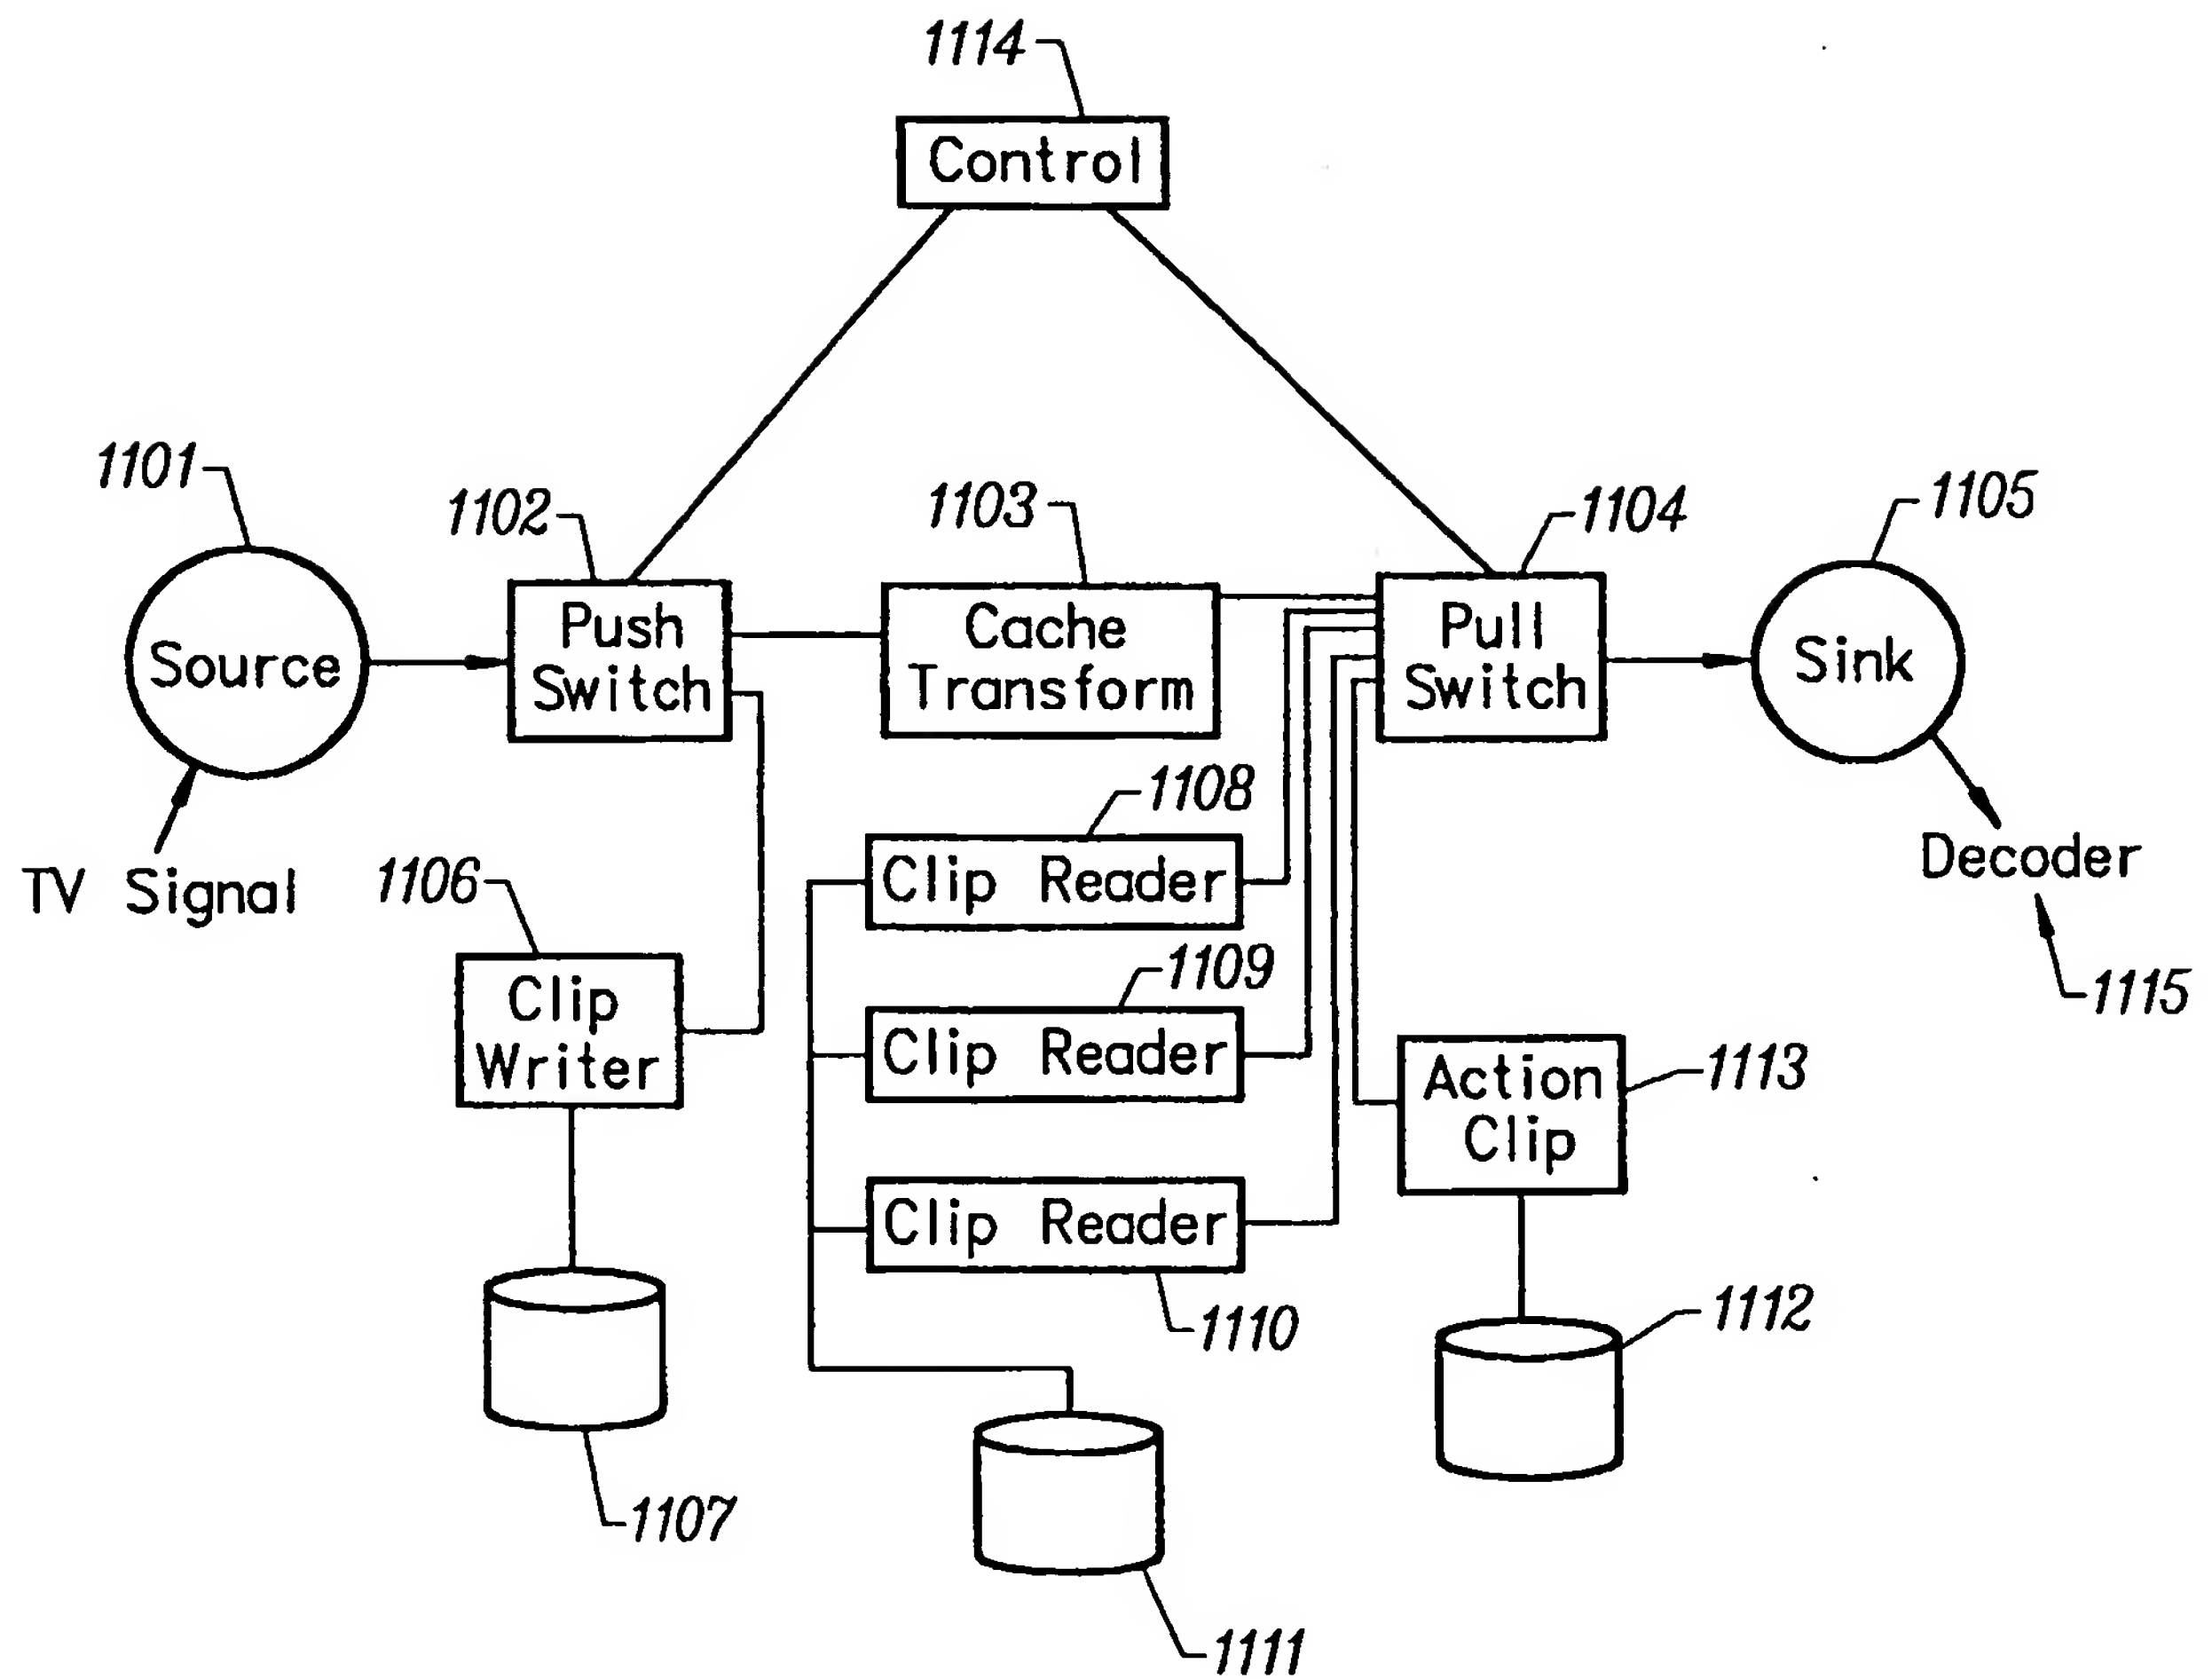


FIG. 11

11/22

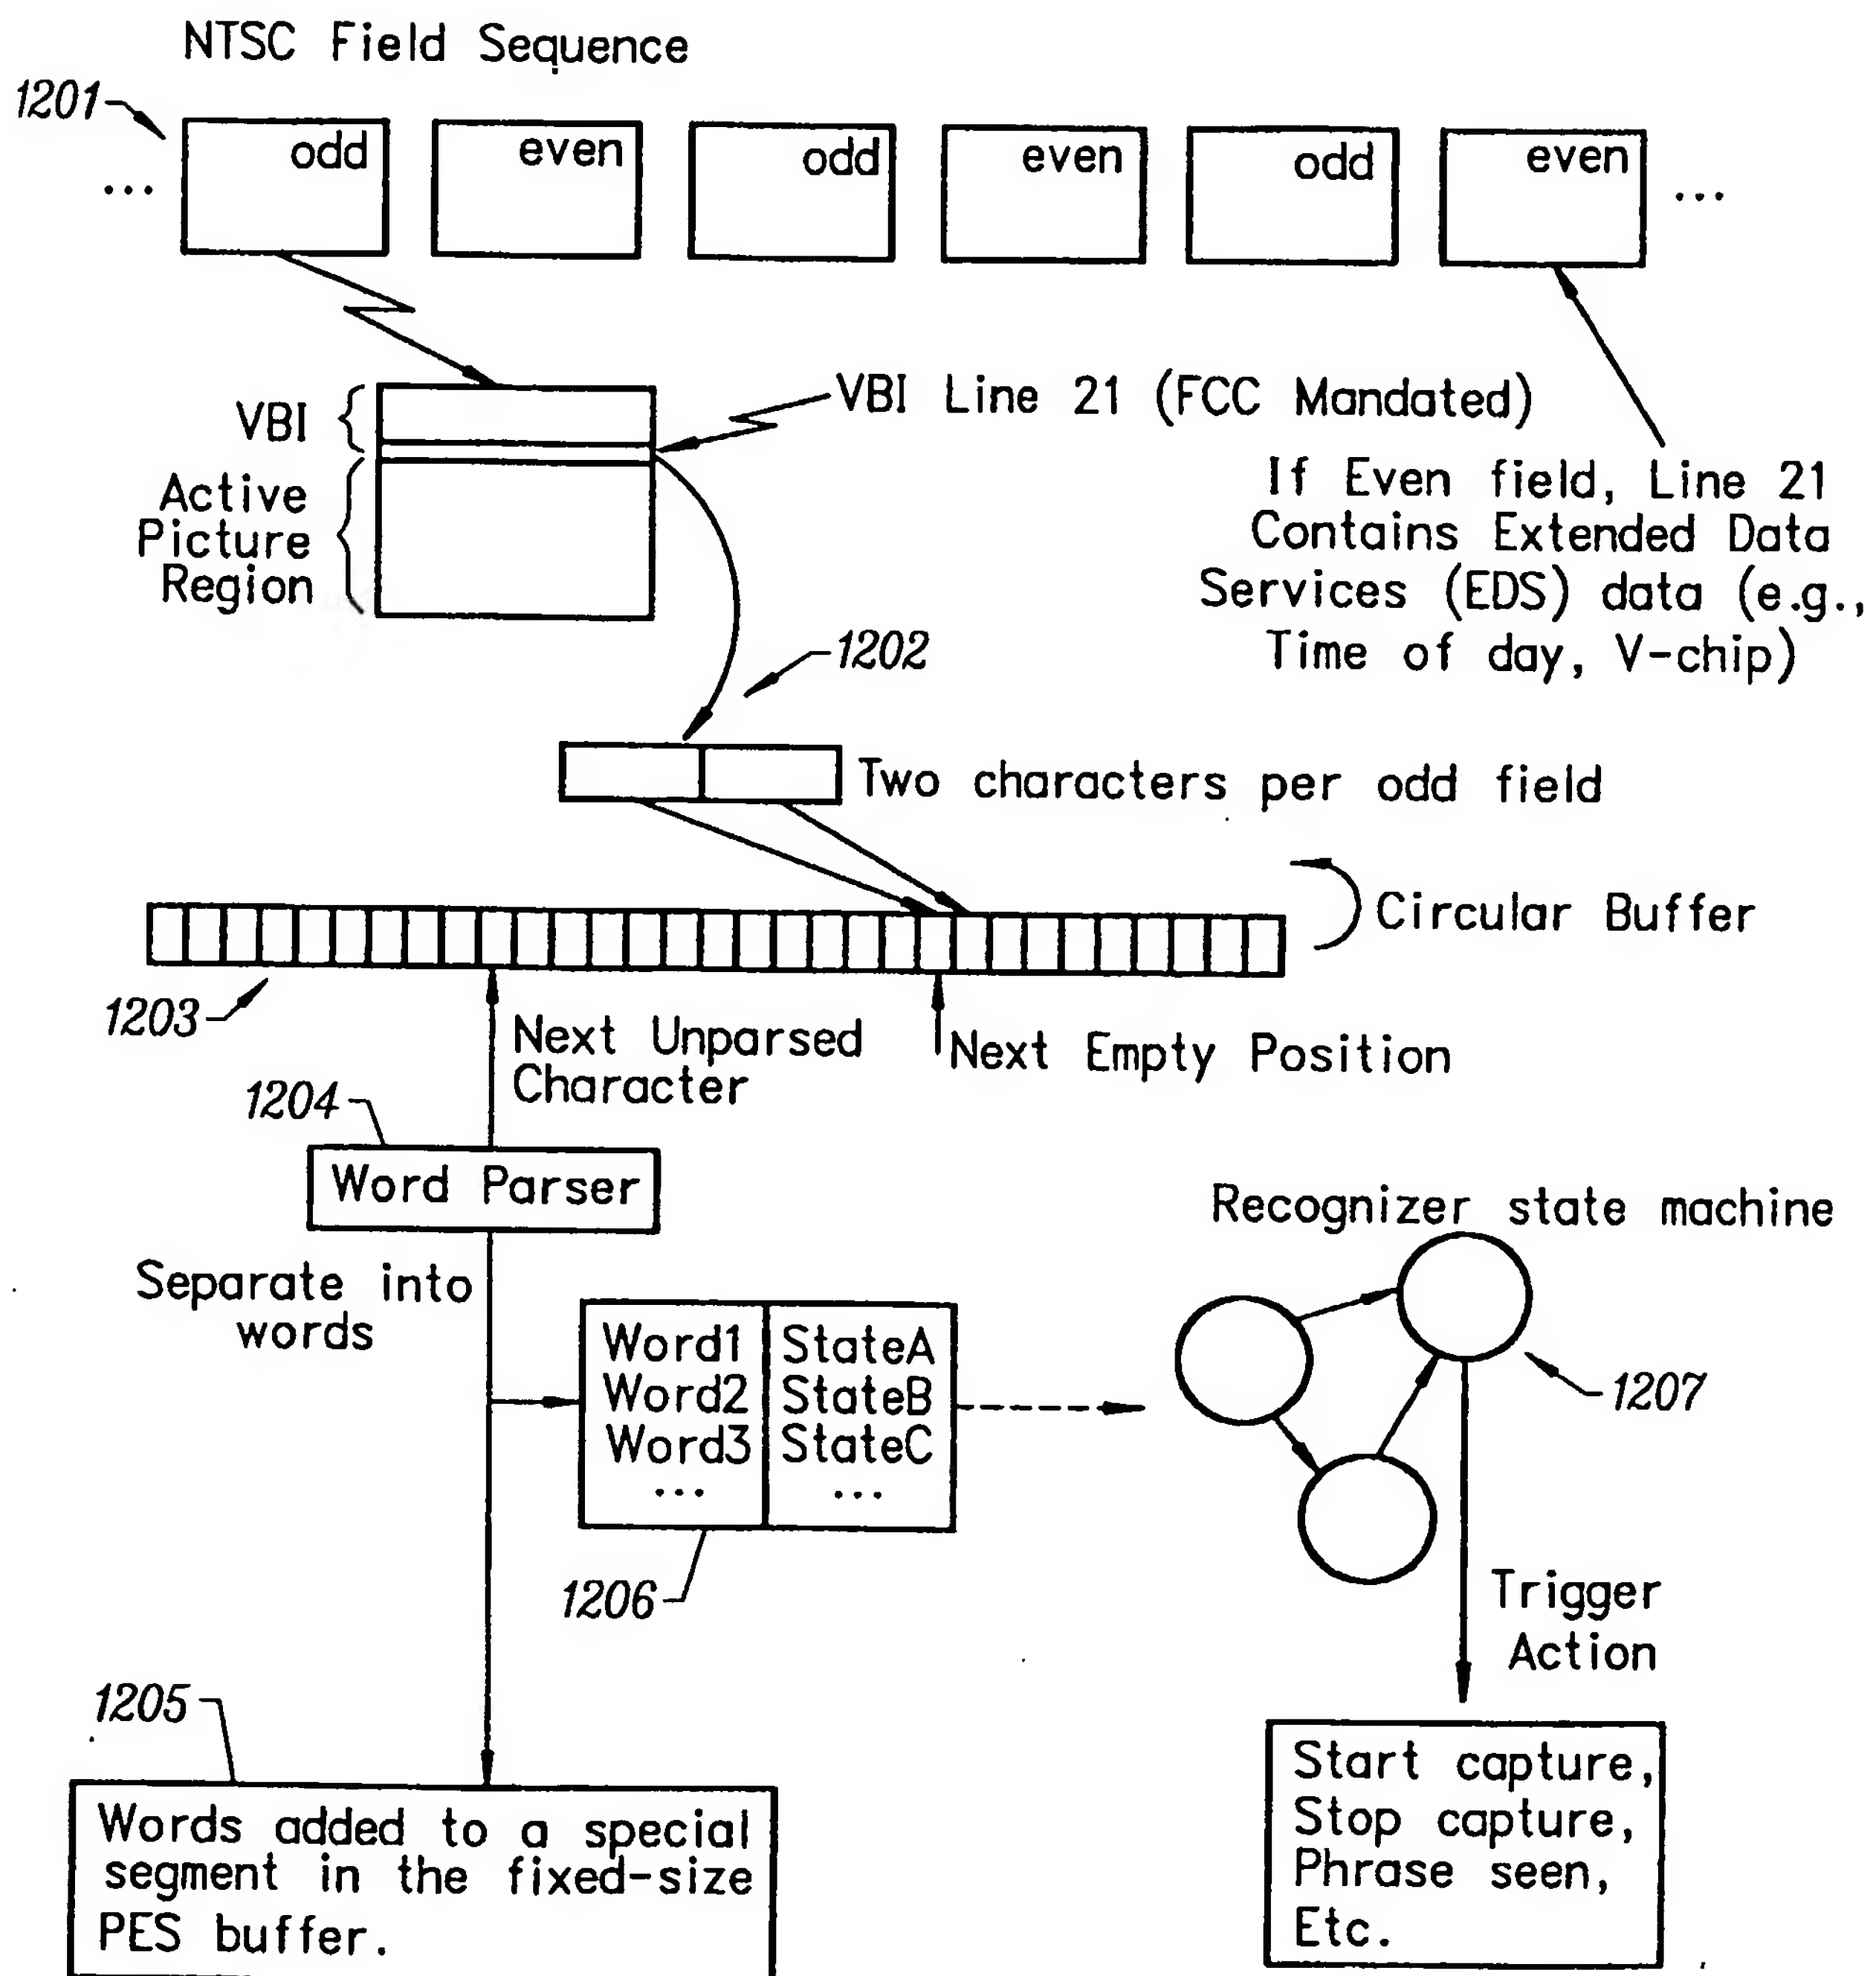




FIG. 12

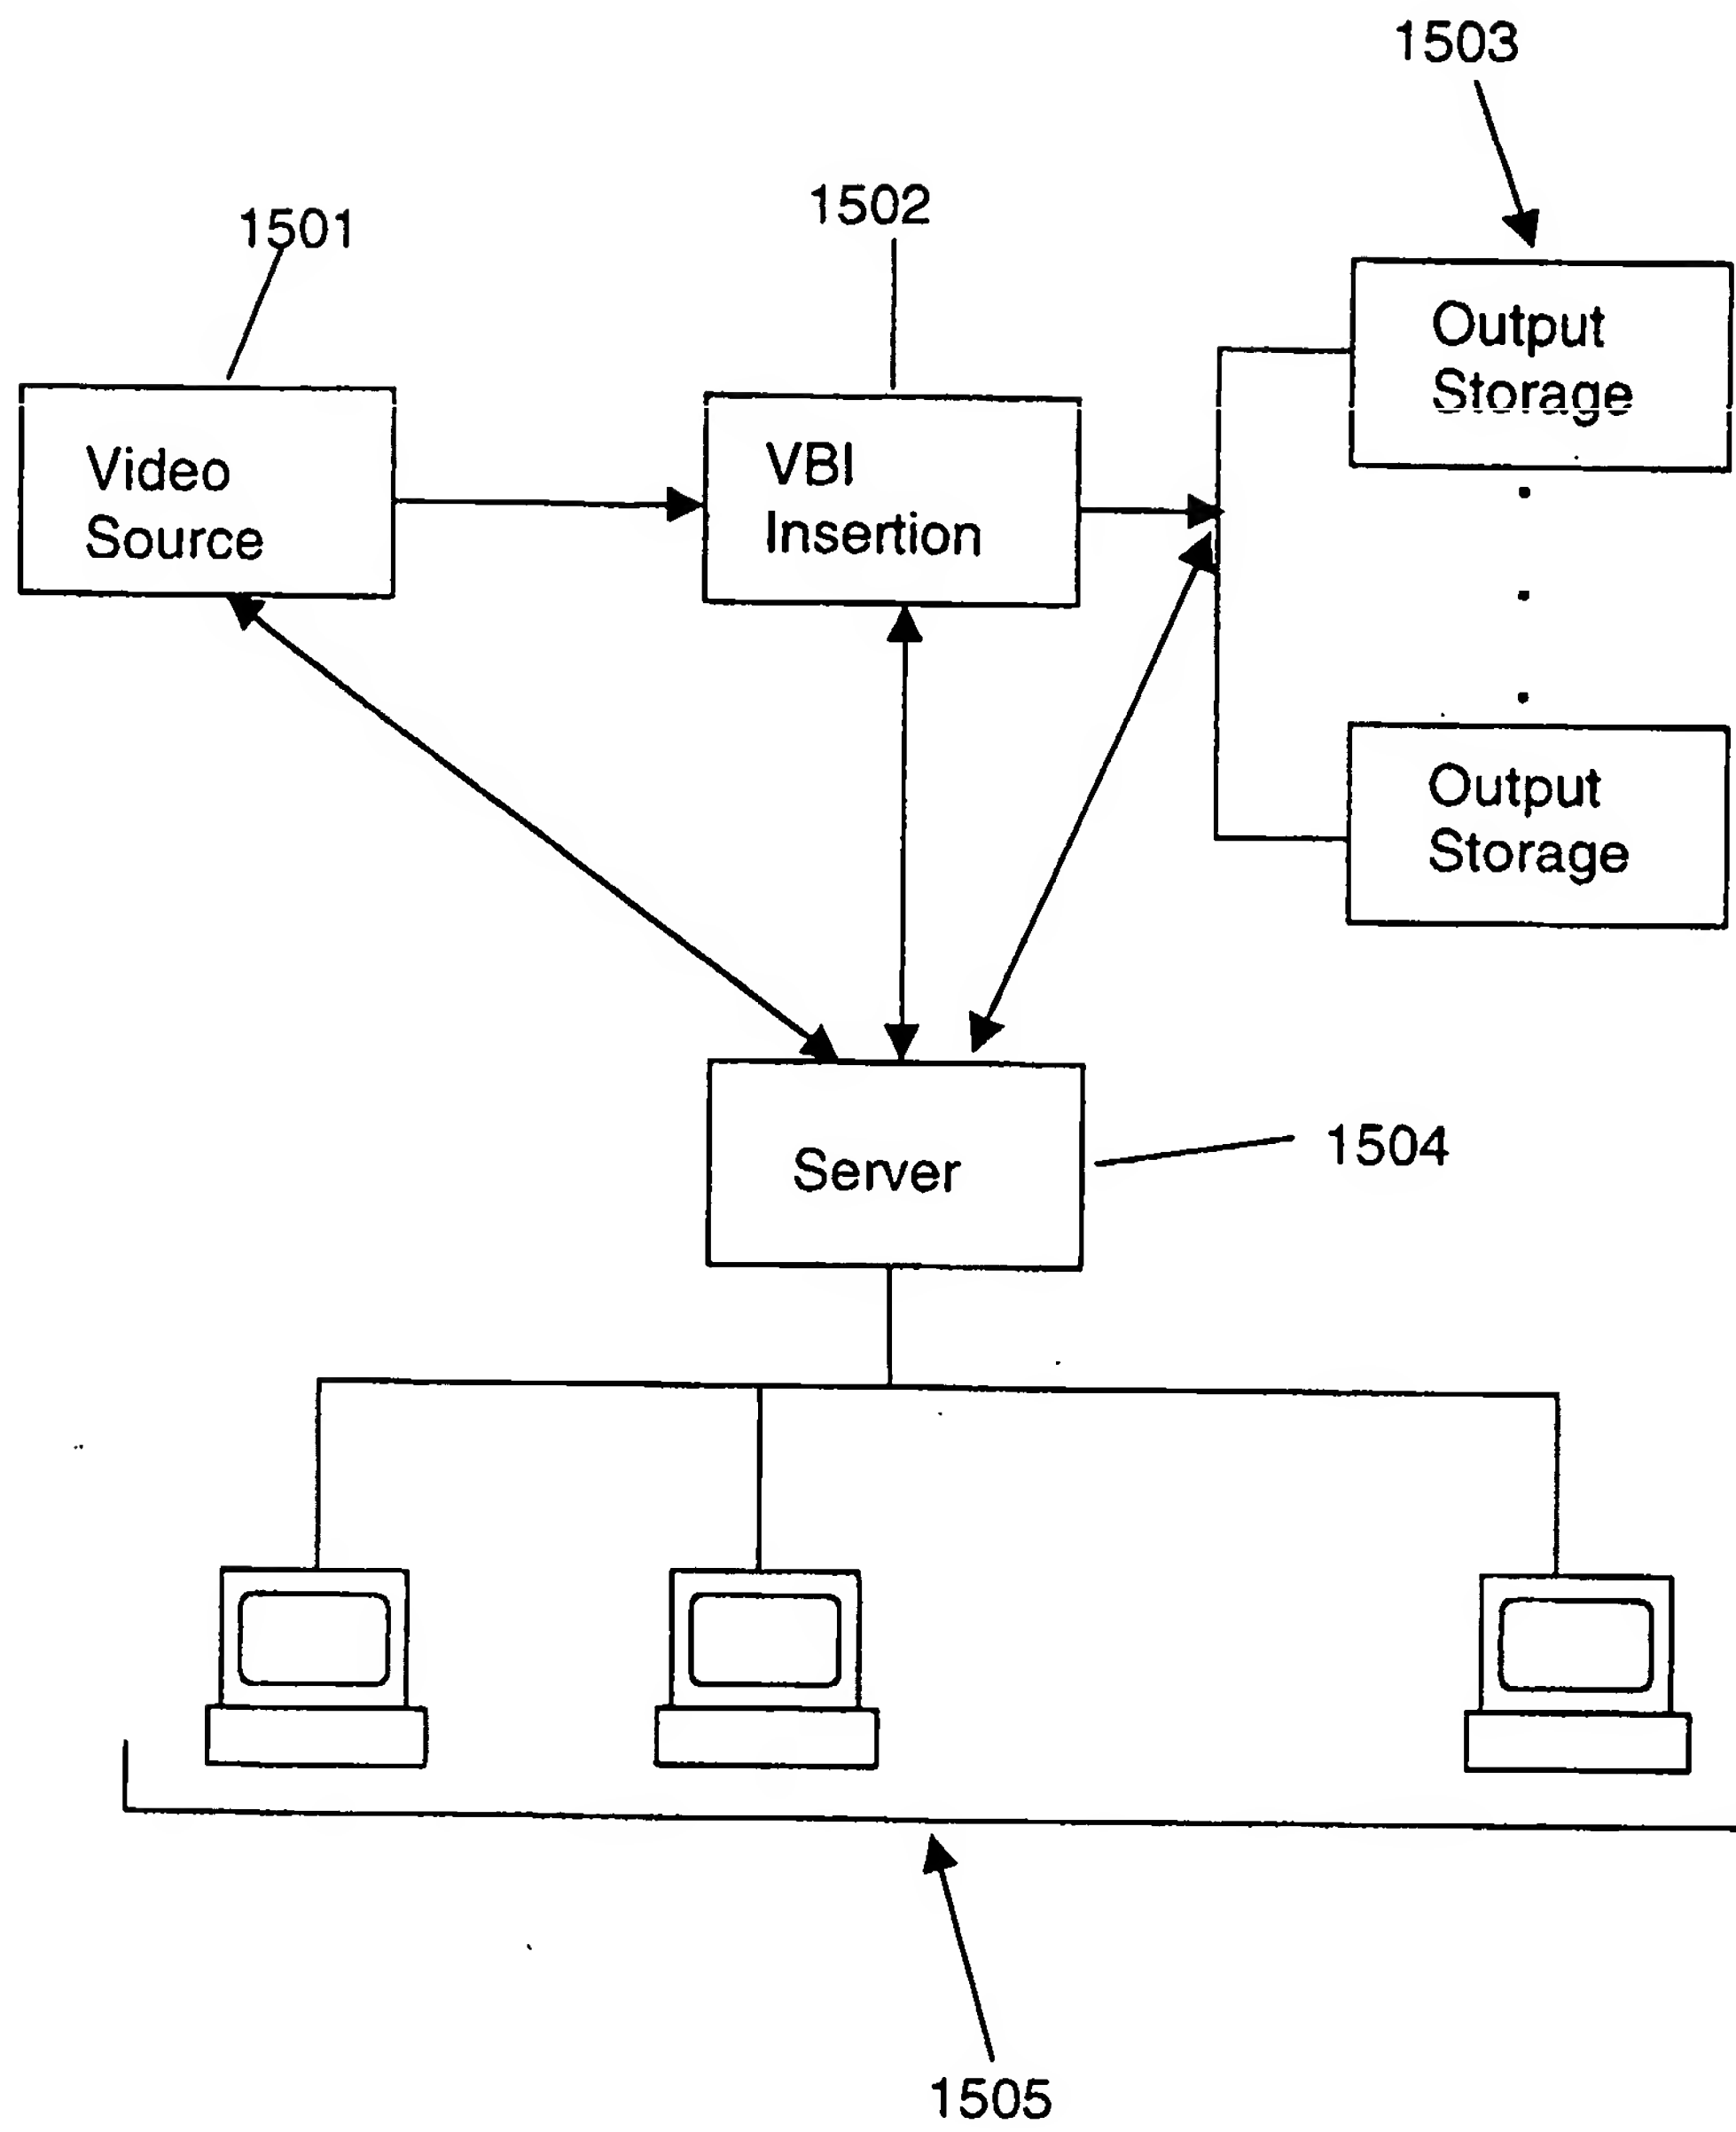
12/22

Fig. 13

13/22

Fig. 14

14/22

Fig. 15

15/22




Fig. 16

16/22




Fig. 17

17/22




Fig. 18

18/22


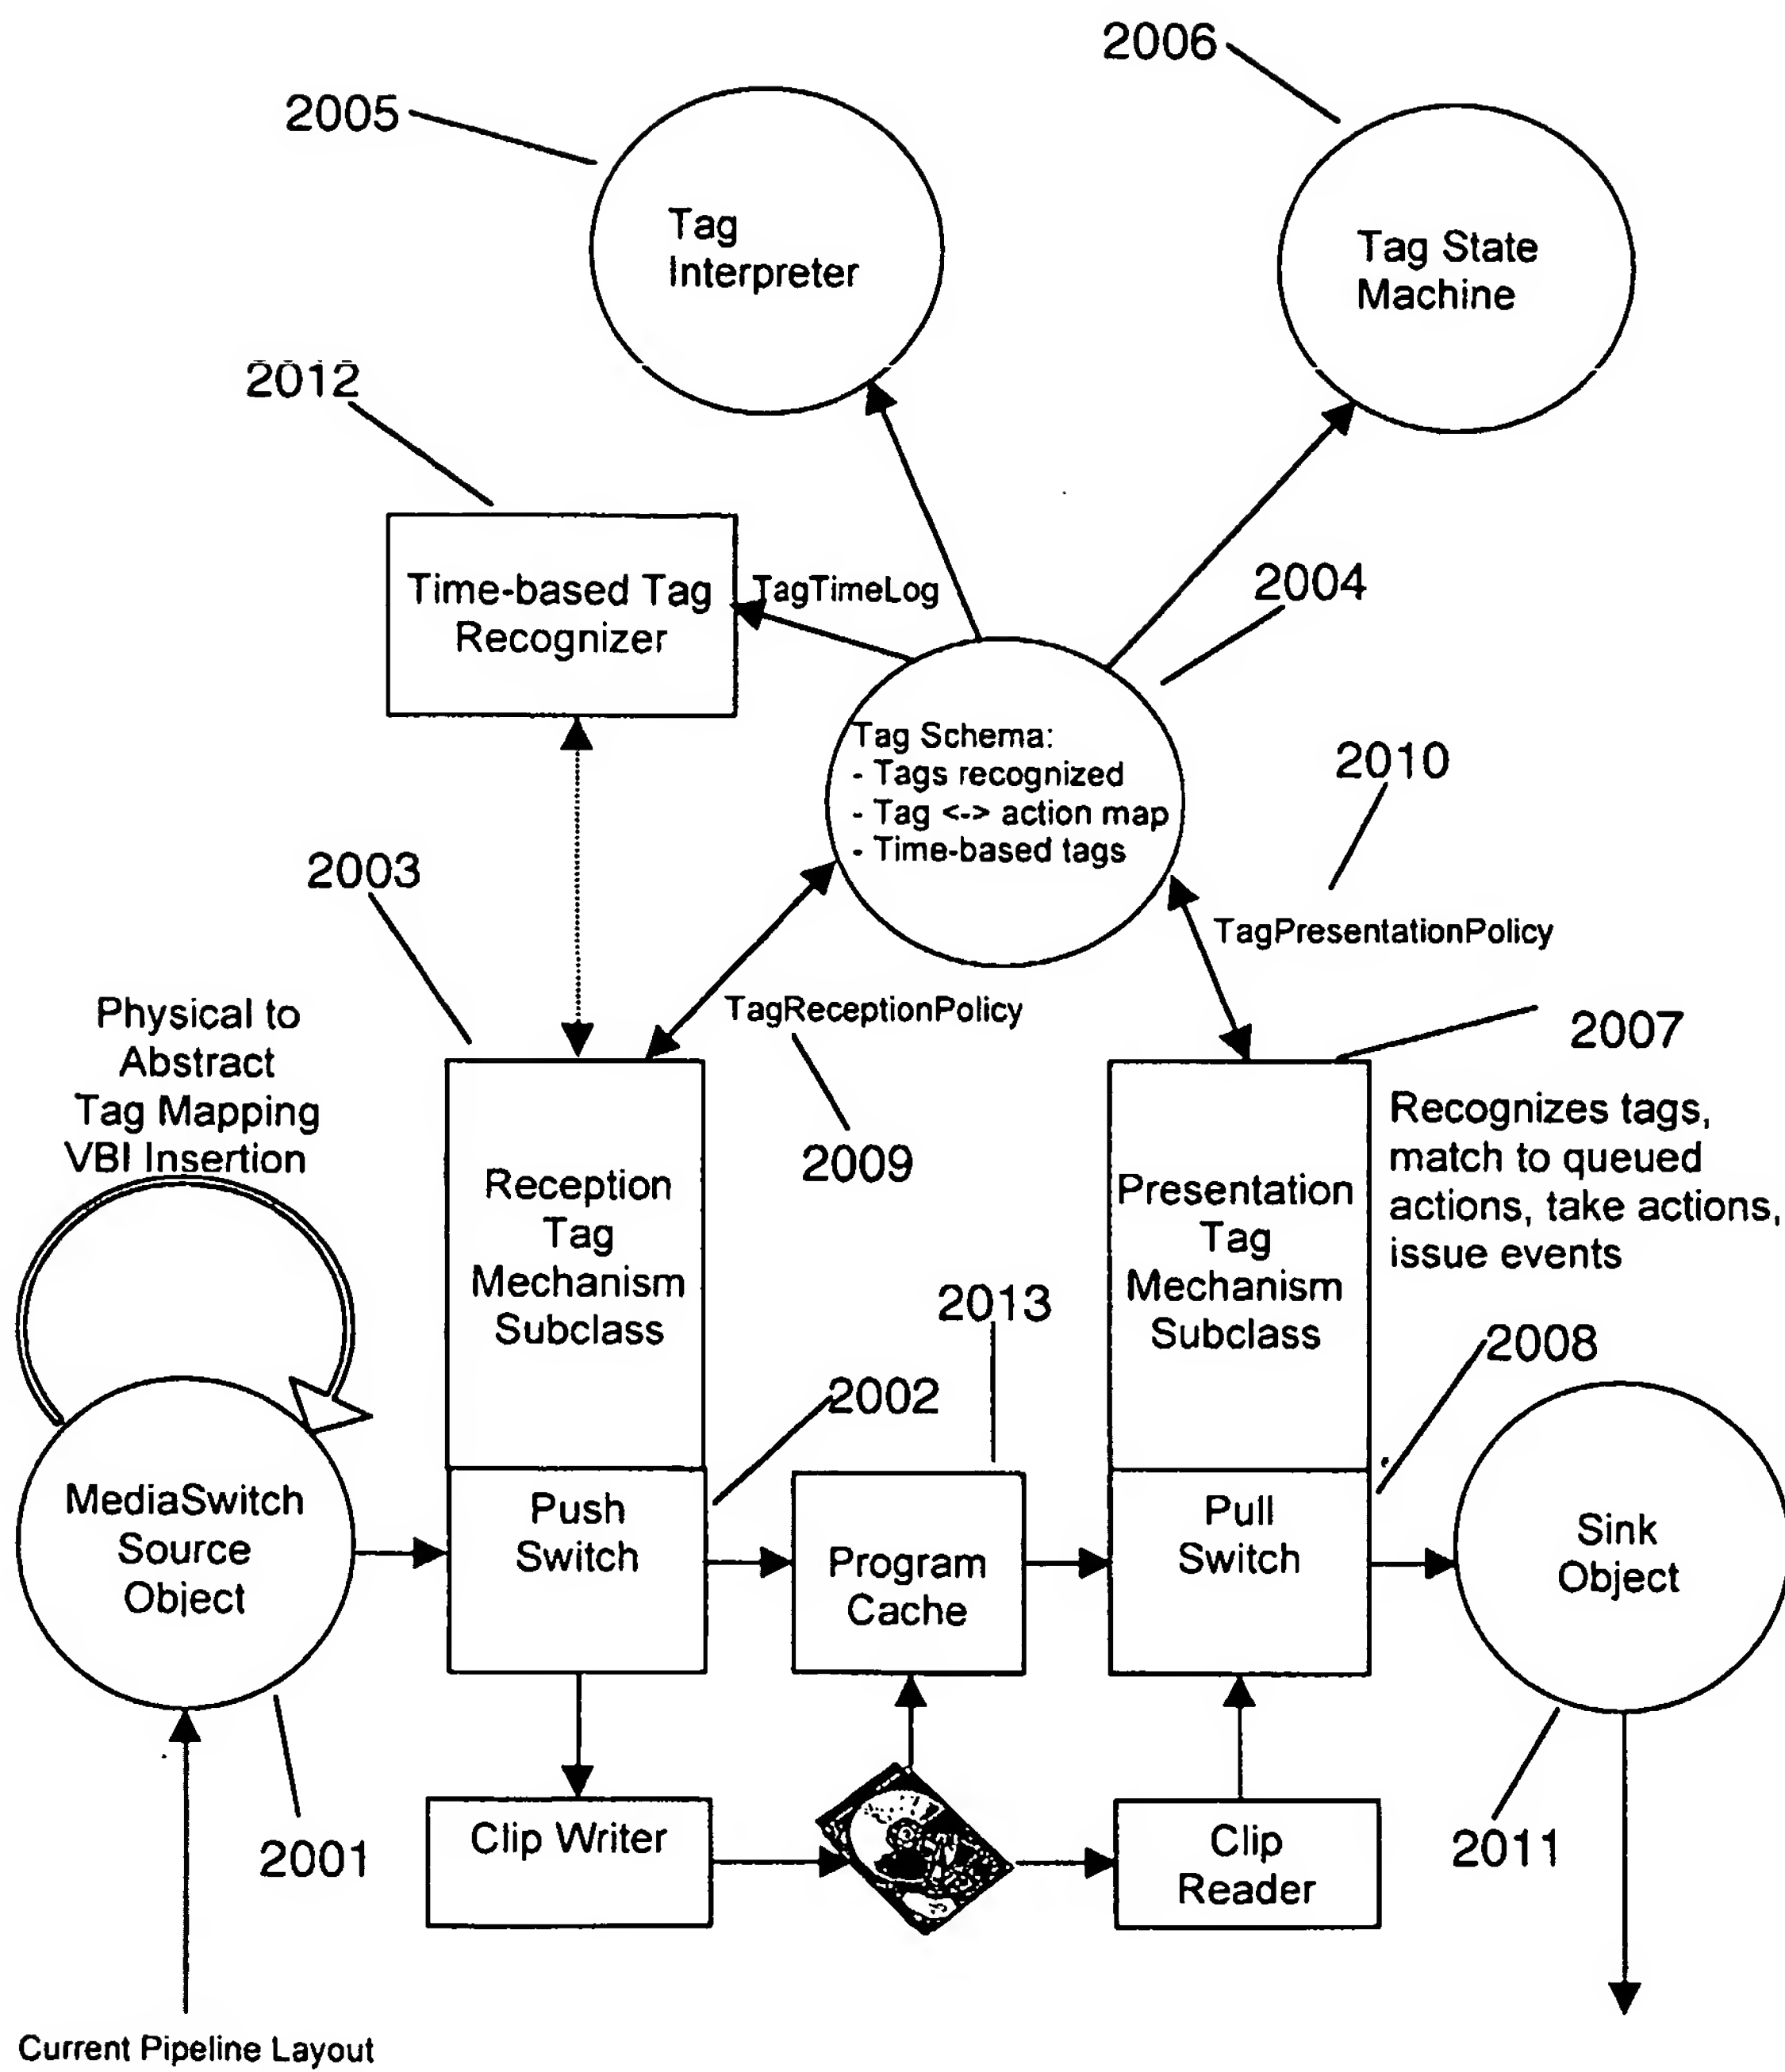
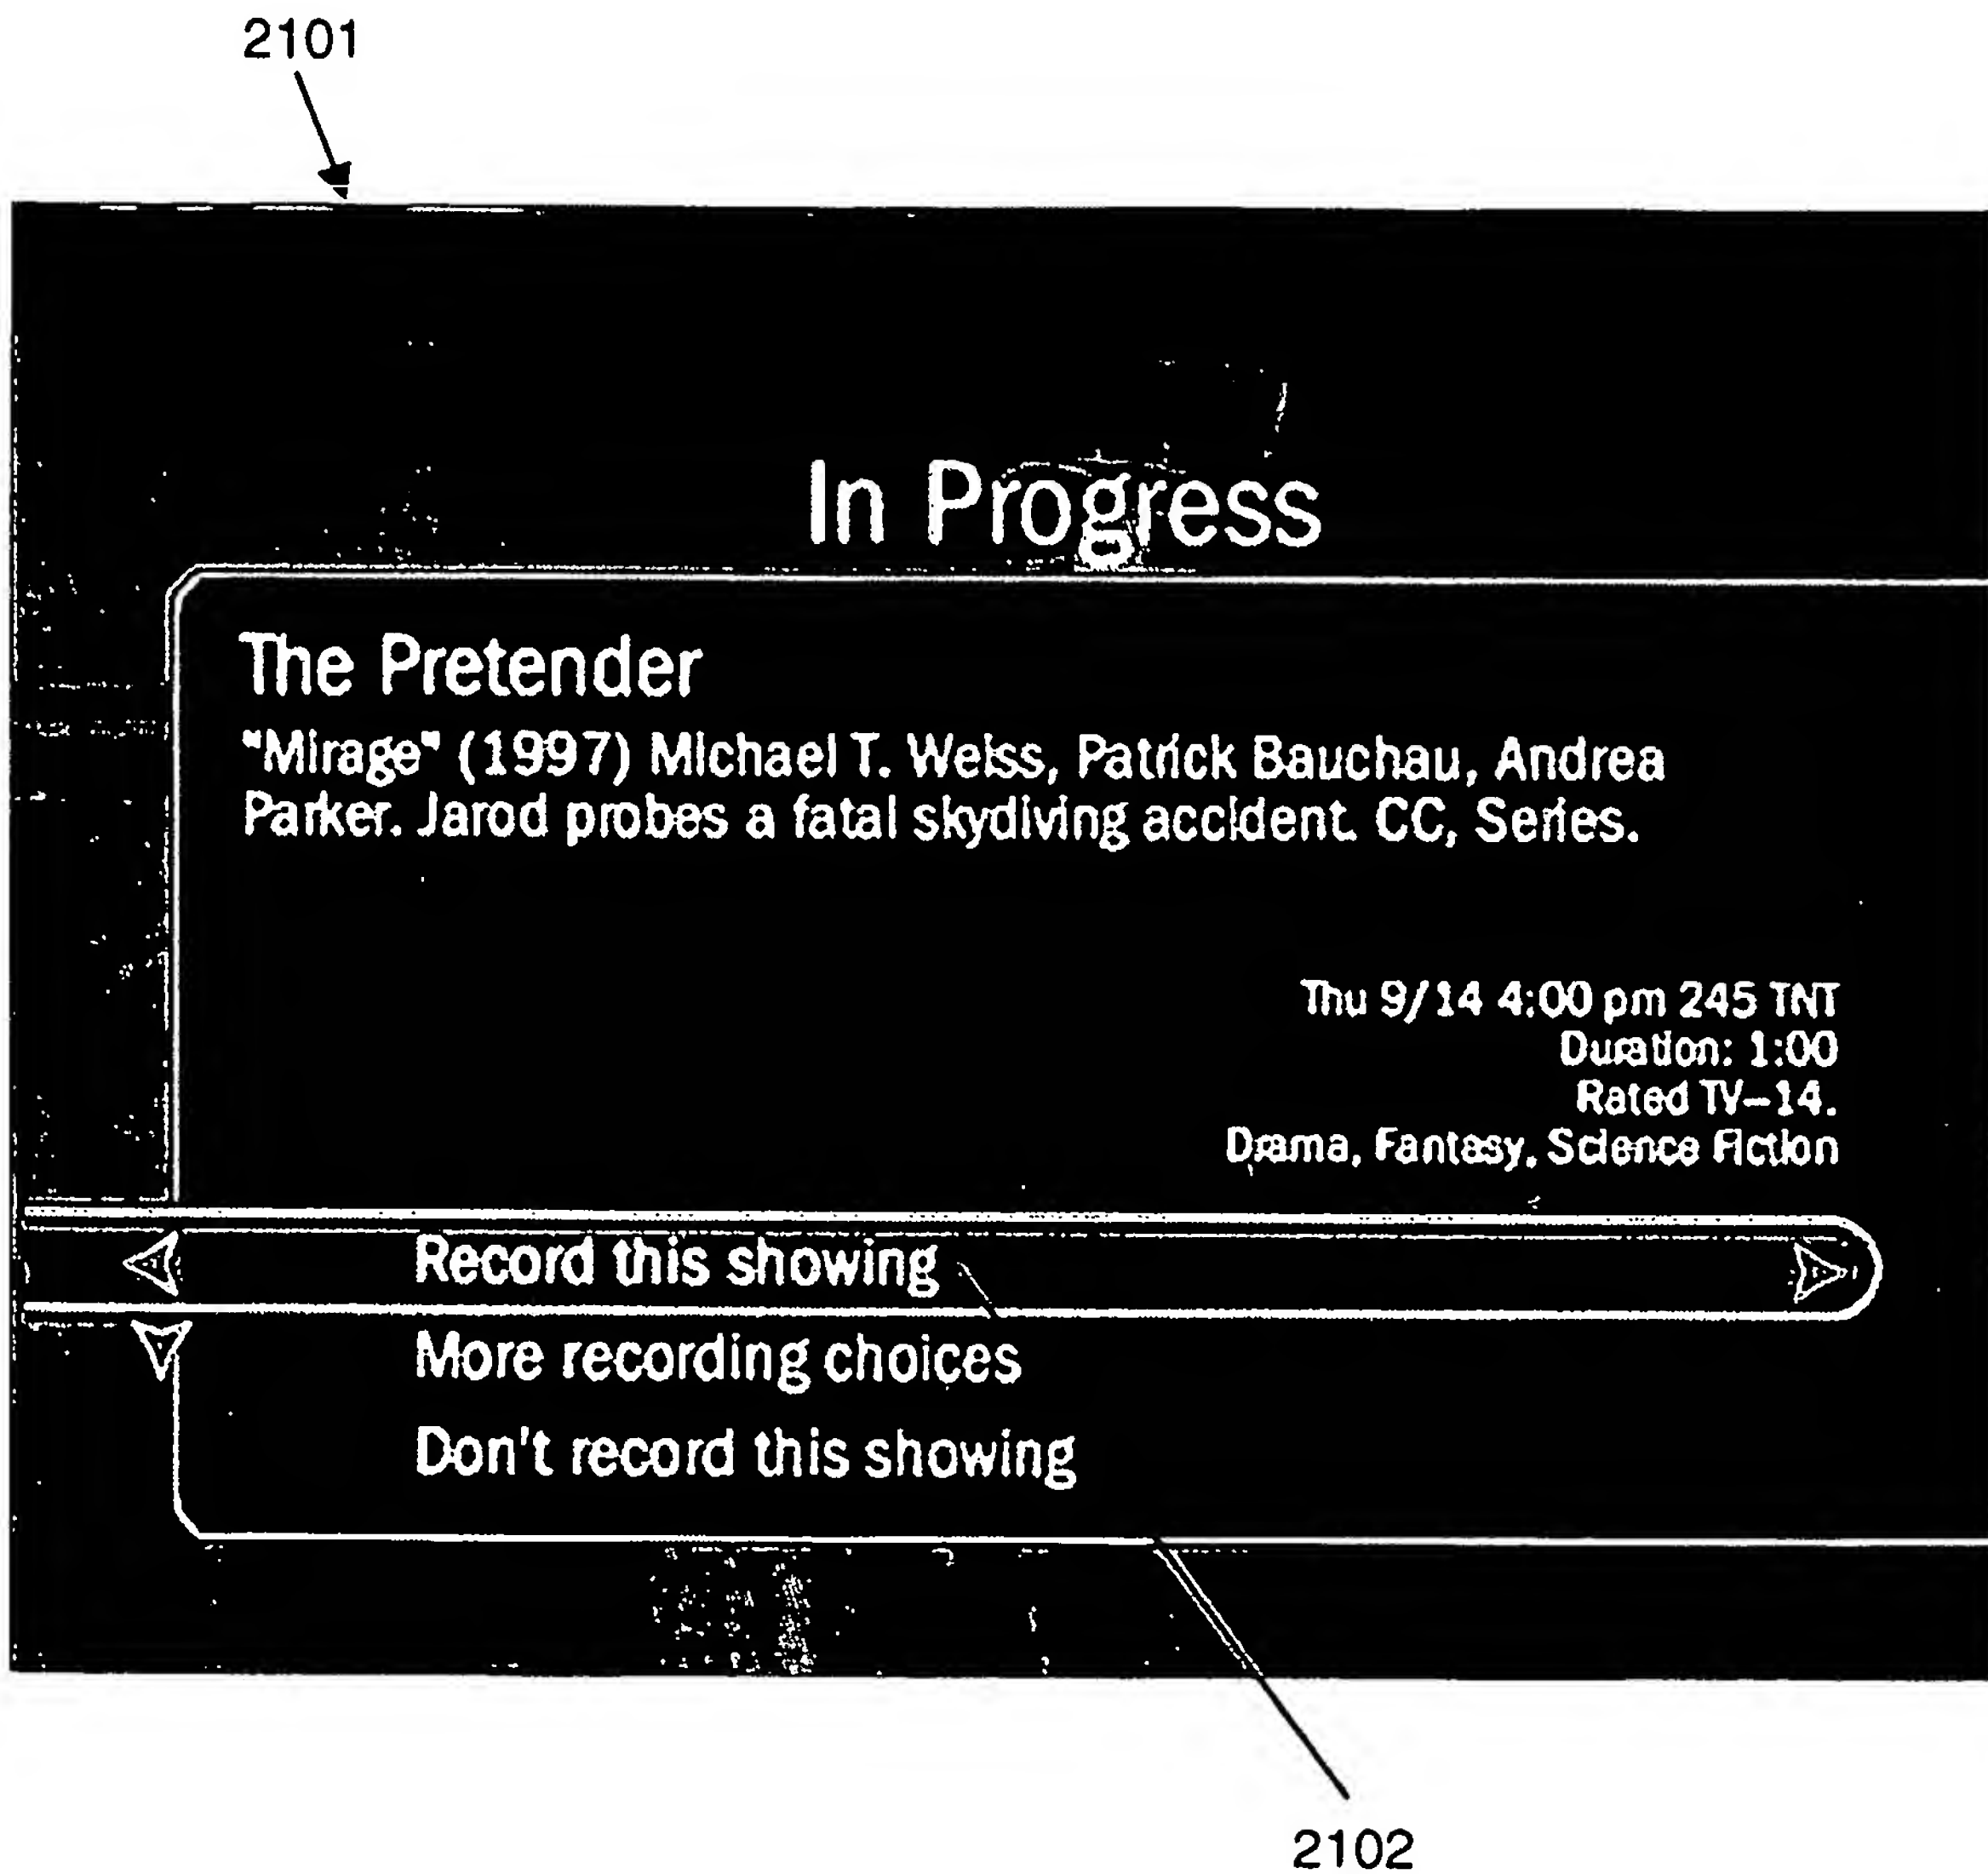


Fig. 19

19/22

Fig. 20

20/22

Fig. 21

21/22




Fig. 22

22/22

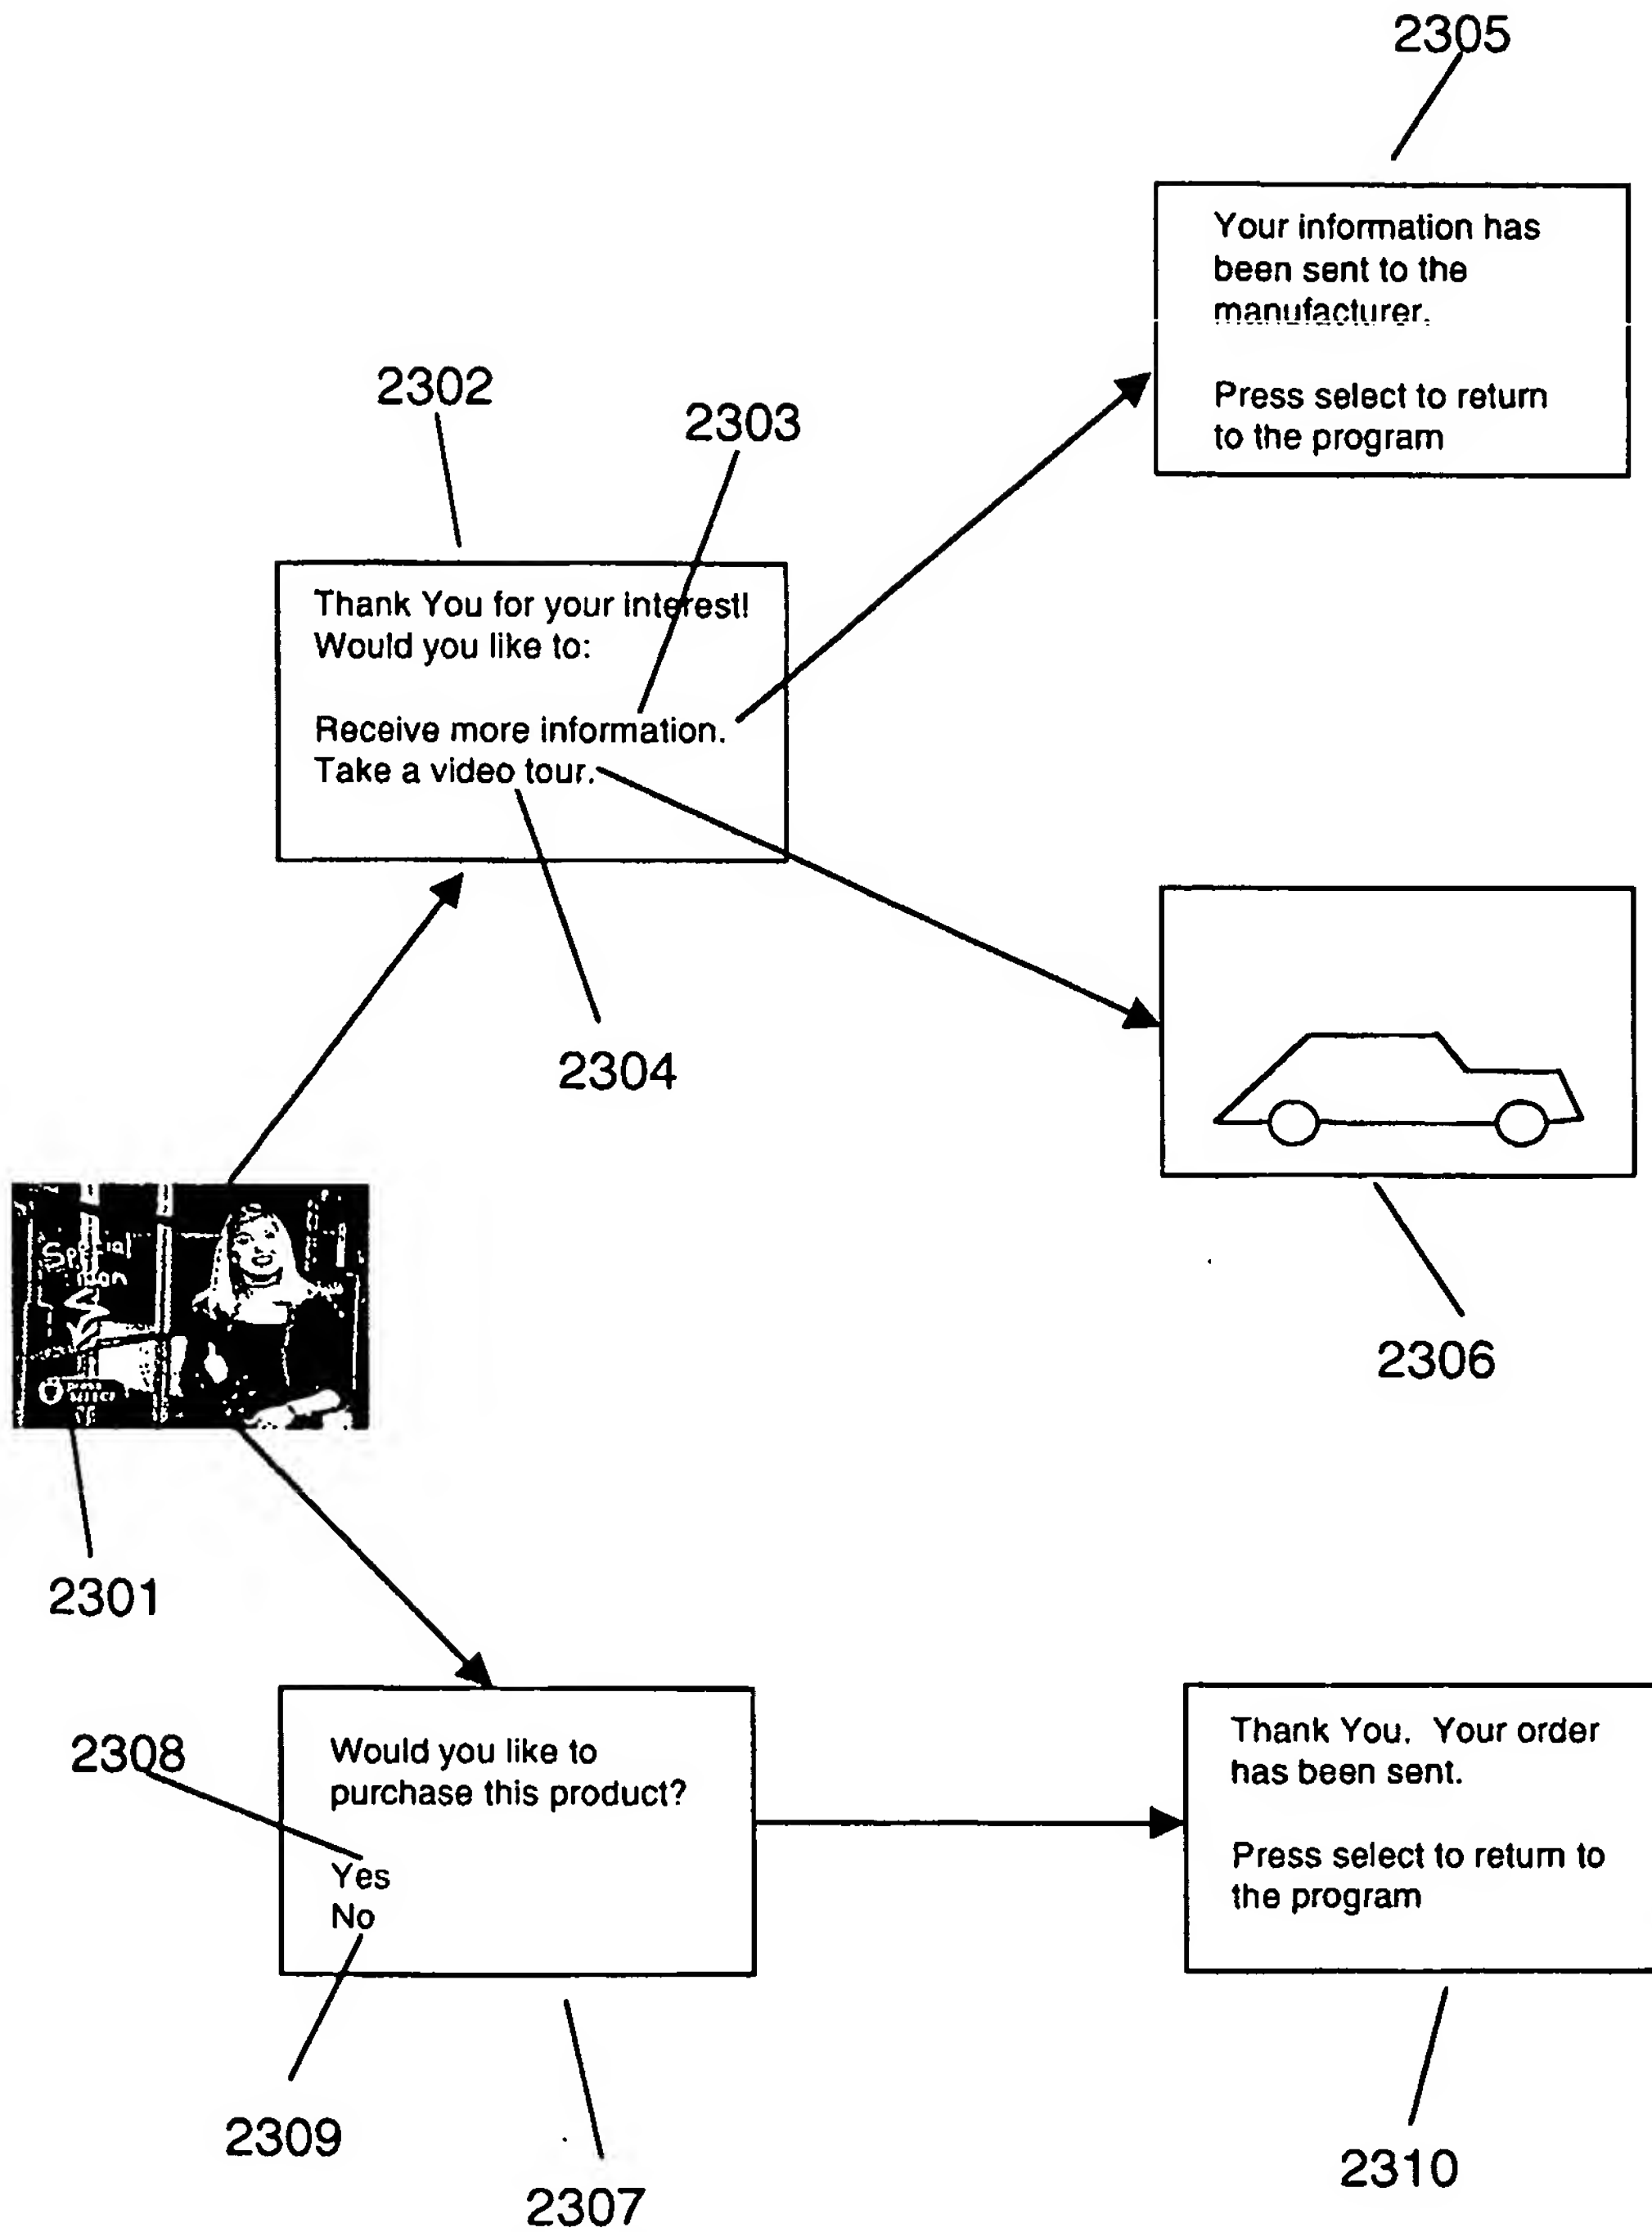


Fig. 23

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04N5/92

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 96 08921 A (ARTHUR D. LITTLE INTERPRISES, INC) 21 March 1996 (1996-03-21) page 9, line 15 -page 27, line 20; figures 1-10	1-4, 28-31
A	DE 44 34 034 A (DEUTSCHE THOMSON-BRANDT GMBH) 28 March 1996 (1996-03-28) the whole document	1-4, 6, 28-31, 33
A	US 5 805 763 A (LAWLER ET AL.) 8 September 1998 (1998-09-08) the whole document	1, 5-8, 10, 11, 15, 17, 28, 32-35, 37, 38, 42, 44

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *&* document member of the same patent family

Date of the actual completion of the international search

9 January 2001

Date of mailing of the international search report

15/01/2001

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Verleye, J

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 233 423 A (JERNIGAN ET AL.) 3 August 1993 (1993-08-03) the whole document -----	18, 45

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9608921 A	21-03-1996	US 5696866 A AU 694568 B AU 3680295 A BR 9508912 A CA 2199485 A EP 0781486 A JP 10507884 T US 5987210 A US 5999688 A	09-12-1997 23-07-1998 29-03-1996 02-06-1998 21-03-1996 02-07-1997 28-07-1998 16-11-1999 07-12-1999
DE 4434034 A	28-03-1996	NONE	
US 5805763 A	08-09-1998	NONE	
US 5233423 A	03-08-1993	NONE	